The Price of Exclusion: SME Set-Asides in Public Procurement

Darcio Genicolo-Martins[a]

[a] *INSPER, São Paulo*

Abstract

Set-asides expand SME access to public procurement by removing rival bidders from the auction. This paper shows that, when the excluded bidders are the price-forming ones, the policy replaces competition with eligibility at high static cost. I study São Paulo's centralized electronic procurement platform, where a legal reinterpretation expanded SME-only tendering into medical and hospital supplies. The platform's reverse auctions record drop-out prices that, under the maintained independent-private-values clock interpretation, reveal type-specific willingness to supply. I use these exits to decompose the set-aside price effect into lost competitive discipline from excluding non-SMEs and a protected-pool offset from the post-policy SME pool. In standardized non-pharmaceutical procurement, the protected pool responds but does not replace the excluded discipline: the full set-aside generates a static welfare loss of 28.9 percent of the open-regime price at $\lambda = 0.30$. Pharmaceutical procurement exhibits larger but more model-sensitive losses. A 10 percent SME price preference, simulated as a static design benchmark, keeps non-SMEs in the auction at near-zero price cost in standardized markets, but delivers less redistribution than full exclusion. The relevant frontier is therefore not SME support versus no support; it runs between exclusionary redistribution and support that preserves the price-forming bidder pool.

Keywords: public procurement, SME set-asides, auctions, welfare

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

JEL MSC: H32, H57, L26, L53, D44

1. Introduction

Set-asides open SME access to public contracts by closing the auction to everyone else. The design is mechanical: protected firms come in, unprotected firms go out. Removing rival bidders weakens the competitive discipline that forms the price. Admitting more SMEs into a now-restricted contest pushes back. The static price and welfare case depends centrally on which force is larger, and on how the protected SME pool reorganizes once the auction is closed to non-SMEs. The empirical literature has documented the cost of SME-favoring rules but rarely separates these two channels, even though policy design lives in their balance.

I show that a set-aside replaces competition with eligibility, and that when the bidders it removes are the price-forming ones the replacement is partial, expensive, and inferior to a price preference as a static design benchmark in standardized markets. The contribution is not to show that SME procurement support can be costly, nor to introduce the distinction between set-asides and preferential instruments; both points are central to the procurement literature already. Marion (2007) and Nakabayashi (2013) document that small-business procurement policies can raise procurement costs. Krasnokutskaya and Seim (2011) show that bid-preference programs affect participation and procurement outcomes, and Athey et al. (2013) compare set-asides and subsidies in an auction framework where instrument design shapes entry, prices, and efficiency. A broader recent literature anchors the welfare-relevance of procurement design across settings (Coviello and Mariniello, 2014; Bosio et al., 2022; Best et al., 2023; Decarolis et al., 2025). This paper instead moves the comparison from average policy costs to price formation. The missing object is not whether SME-favoring rules affect procurement outcomes, but which bidder pool supplies the price-forming order statistic once the rule excludes non-SMEs.

I make three contributions. First, I decompose the price effect of an SME set-aside into a lost-discipline component (non-SMEs removed from the price-forming pool, SME pool held fixed) and a protected-pool offset (post-policy SME pool replacing the pre-policy one). Second, I implement this decomposition in a reverse-auction setting where losing bidders' drop-out prices reveal willingness to supply under the maintained IPV clock interpretation, so the price-forming order statistic can be simulated directly. Third, I recast policy design: the relevant frontier is not SME support versus no SME support, but exclusionary redistribution, which buys SME access by removing rival bidders, versus competition-preserving redistribution, which tilts allocation while keeping non-SMEs inside the auction. The 10 percent price preference enters as a static design benchmark, not as a full implementation forecast: it preserves the price-forming pool, at the cost of delivering less redistribution than a full set-aside.

The setting is São Paulo's centralized electronic procurement platform, the Bolsa Eletrônica de Compras (BEC). In late 2017 and early 2018, a legal reinterpretation moved the state's SME-only eligibility threshold from the total value of a purchase notice to the item level. The change mechanically pushed medical and hospital supplies (Group 65) into SME-only tendering, because such purchases combine many small items inside large notices. The empirical cutoff is March 2018, when Group-65 public buyers began mass adoption of the new functionality. The trigger was legal and administrative, not a medical-supply policy or a technological shock.

The first-stage facts fit the institutional account. After the cutoff, SME-only adoption in Group 65 rises sharply, non-SME participation falls, SME participation roughly doubles, and winning prices rise. A compact difference-in-differences benchmark against 76 never-treated product groups places the associated price movement at 10–11 percent in the structural window. The DiD is not asked to carry the structural magnitude; its role is to verify the timing, sign, and approximate scale of the predicted footprint. Under explicit auction assumptions, the structural decomposition then identifies how price formation moves when the policy removes non-SMEs and reshuffles the active SME pool.

The auction format makes the decomposition feasible. BEC's main format is the electronic *Pregão*, an English-reverse auction. Under the maintained independent-private-values clock interpretation, losing bidders' drop-out prices reveal willingness to supply, which I interpret as costs in the structural model (Haile and Tamer, 2003). This is the identifying advantage of the setting and, equally, its main structural restriction. I use drop-outs to recover type-specific willingness-to-supply distributions (cost distributions under the maintained model), correct for auction-level scale heterogeneity in the spirit of Krasnokutskaya (2011), and simulate counterfactual auctions under observed equilibrium entry. The exercise is neither a full entry equilibrium nor a reduced-form treatment-effect estimate: entry rates enter as primitives, and the post-policy SME willingness-to-supply distribution is estimated directly from post-policy exits rather than derived from a formal selection model.

The decomposition compares three counterfactual auctions on the same product cells. The benchmark is the open auction with the pre-policy bidder pool, S_1. A second counterfactual, S_2, removes non-SMEs while holding the pre-policy SME pool fixed; it isolates the cost of weakening the price-forming order statistic. A third, S_3, additionally replaces the fixed SME pool with the active post-policy SME pool observed under the rule, picking up both additional SME participation and changes in the composition of active SME bidders. The move from S_1 to S_2 measures the lost competitive discipline from excluding non-SMEs; the move from S_2 to S_3 measures the offset that the protected pool generates after the policy; the realized set-aside price effect, S_3 relative to S_1, is the sum of these two channels. The decomposition turns a set-aside price effect into a design question: how much is mechanical exclusion, and how much is the protected pool's response?

The main result is that the protected pool responds, but exclusion dominates. In non-pharmaceutical supplies, removing non-SMEs while holding the pre-policy SME pool fixed raises simulated prices by 0.371 of the reference price; the observed post-policy SME pool offsets that shock by -0.144, leaving a net effect of 0.227, with the exclusion channel accounting for 72.0 percent of the two components in absolute magnitude. The protected pool expands and changes composition, but it does not recreate the competitive discipline supplied by excluded non-SMEs. In pharmaceutical procurement the same qualitative pattern appears at larger magnitudes, but composition turnover makes the welfare ranking conditional. I report pharmaceuticals as a boundary case rather than as a second headline.

The result reframes policy design. A full set-aside supports SMEs by removing their rivals; a price preference does so while keeping non-SMEs in the auction. I simulate that benchmark under fixed entry and recovered willingness-to-supply primitives. In thick, standardized non-pharmaceutical markets, the preference has near-zero price and welfare cost while the full set-aside generates a substantial loss. The preference is not a free version of the set-aside; it delivers smaller redistribution. The comparison is therefore a frontier, not a ranking: a planner should choose the set-aside only when the additional SME surplus is worth the allocative and fiscal cost of removing non-SMEs.

The strongest policy result is in thick standardized markets. In thinner pharmaceutical markets, the ranking depends on how the post-policy SME willingness-to-supply distribution is modeled. The non-pharmaceutical ranking is stable across the main and strict-invariance readings; the pharmaceutical ranking is not. I treat the gap as a scope condition rather than a second headline: bidder exclusion is hardest to defend where the protected pool is thick enough for a preference rule to tilt allocation without removing competitive discipline, and most defensible, if at all, where the protected pool is thin, the good is heterogeneous, and the planner places high value on SME surplus.

A welfare comparison in the spirit of Saez and Stantcheva (2016) maps the two channels into transfers to SMEs, allocative waste, and the marginal cost of public funds. The lesson is direct: when the bidders a set-aside removes are the ones forming the price, the rule trades competitive discipline for an incomplete protected-pool response and earns the difference as a welfare loss. Full set-asides remain defensible when the planner assigns enough weight to SME surplus, market access, or dynamic capacity gains beyond what the protected pool itself delivers. The natural benchmark for SME procurement support is not no support, but support that keeps rival bidders inside the auction.

2. Institutional Setting, Data, and First-Stage Facts

The empirical episode is a legal and administrative change that shifted eligibility for Group 65 procurement items. Its immediate effect was to change the bidder pool: SME-only take-up rises, non-SME participation falls, SME participation rises, and prices move up around the March 2018 mass-take-up date. These facts establish the timing and direction of the first stage. They do not, by themselves, identify the structural price mechanism; that requires the auction decomposition developed in Section 3. This section documents the legal shock, the data and structural sample, the first-stage participation and price facts, and the identification architecture that connects them to the model.

2.1. Legal shock

The treatment is not the existence of Brazil's SME procurement statute. It is the expansion of SME-only eligibility induced by a change in how São Paulo applied the R$ 80,000 threshold. Brazil's federal SME statute, LC 123/2006, requires exclusive SME tenders for procurement items valued at R$ 80,000 or less, and LC 147/2014 made the SME-only rule mandatory rather than optional. For most goods on São Paulo's BEC platform, that rule already operated before the period studied here.

Group 65—medical, dental, and hospital equipment and supplies—was different for a legal reason. The prevailing interpretation applied the R$ 80,000 threshold to the total value of the purchase notice, not to each item or lot. Because medical-supply notices often bundle many individually small items inside large purchase orders, the total-notice interpretation kept many individually eligible items outside the SME-only default. The item-level interpretation moved those items into SME-only eligibility.

The institutional sequence supports March 2018 as an operational cutoff. BEC enabled SME-only purchase-order functionality in COMUNICADO BEC nº 02/2017 on 18 July 2017 and mixed-exclusivity notices in COMUNICADO BEC nº 03/2017 on 20 November 2017. PGE-SP then issued Parecer Sub-G Cons. nº 151/2017 on 11 December 2017, holding that the R$ 80,000 threshold applies item-by-item regardless of the total value of the purchase offer. The empirical cutoff is March 2018, when Group 65 public buyer units began mass take-up of the functionality. I use this operational adoption date rather than any single legal document as the cutoff. TCE-SP aligned with the new interpretation on 16 May 2018, after the empirical cutoff, ratifying rather than initiating the operational sequence.

The institutional claim is narrow. The trigger was legal and administrative, not a Group 65 market intervention, and it did not speak to medical-supply prices, supplier complaints, or production technology. It was a general reading of procurement law that interacted with the way Group 65 notices were written. The cutoff shifts admissibility and the active bidder pool; it does not by itself identify a full entry model or rule out composition changes among SMEs. Online Appendix OA-A reports the institutional timeline and platform sample details.

2.2. Data and auction format

The data are administrative microdata from BEC, São Paulo's centralized electronic procurement platform; the Brazilian institutional setting has been used to study procurement effects on firm dynamics (Ferraz et al., 2016). The raw extract defines the platform universe used in the paper: 3.7 million bid-level observations across 82,569 items, 832,984 purchase orders, and 1,344 public buyer units. Group 65 accounts for roughly 27 percent of platform transactions in the window used here and is the treated group. The reduced-form benchmark uses this broader panel to compare Group 65 with 76 never-treated product groups.

The structural exercise uses Group 65 Pregão auctions in a symmetric 18-month window on each side of the March 2018 cutoff, from September 2016 through August 2019. The Pregão format is central because it is an electronic English-reverse auction: the auction opens at a buyer reference price, firms submit lower offers, and firms exit as the clock falls. Under the maintained independent-private-values clock interpretation, losing drop-out prices are willingness-to-supply observations, interpreted as costs in the structural model.

For that reason, the structural sample restricts attention to Pregão auctions with at least 2 active firms and normalized bids b/p_t^{ref}, denoted c_ϵ, in $(0, 3]$. The resulting sample contains 297,967 firm-auction observations across 97,993 auctions, split almost evenly between the pre-period (48,740 auctions) and post-period (49,253 auctions).

Each observation is classified by bidder type and product class. SME status is taken from the BEC administrative registry and validated against historical firm bidding records. I split Group 65 into non-pharmaceutical supplies and a pharmaceutical class covering CMED-regulated medicines and biologicals. This split is not a second treatment. It is a market-thickness distinction: the non-pharmaceutical segment is closer to a thick standardized market, while pharmaceutical procurement has fewer SME bidders and more within-class heterogeneity. The main policy claim is strongest in the former; the latter is used as a boundary case.

2.3. First-stage facts

The participation margins are the key first stage for the decomposition because the mechanism operates through the bidder pool. Four facts matter. SME-only adoption rises in Group 65 after the cutoff; state-level adherence reaches 43 percent rather than full coverage. Non-SME participation falls. SME participation rises. Winning prices move up in the same period. The participation responses are the moments most useful for the structural exercise because they do not pass through the reference-price normalization; the price movement validates direction and approximate scale, with magnitude estimated structurally in Section 4.

Table 1 reports the participation margins in the structural sample. In non-pharmaceutical Pregão auctions, the average number of SME bidders rises from 0.94 to 1.87, while non-SME participation falls from 2.68 to 1.50. In pharmaceutical auctions, SME participation rises from 0.55 to 1.22, while non-SME participation falls from 2.61 to 1.66. Because take-up is incomplete, post-period non-SME counts remain positive; the first stage is a sharp change in the bidder pool, not universal exclusion.

Table 1: First-stage participation facts in Group 65 Pregão auctions

	Non-pharma	Pharma
Pre: SME bidders per auction	0.94	0.55
Pre: non-SME bidders per auction	2.68	2.61
Post: SME bidders per auction	1.87	1.22
Post: non-SME bidders per auction	1.50	1.66
Change in SME bidders	+0.93	+0.67
Change in non-SME bidders	-1.18	-0.95

Sample: Group 65 Pregão auctions in the 18-month window on each side of the March 2018 cutoff, after the structural filters $c_\epsilon \in (0, 3]$ and at least 2 active firms. Counts are average distinct bidders per auction by type and class. Post-period non-SME counts remain positive because the empirical SME-only take-up rate is 43 percent, not 100 percent.

The price-forming order statistic moves in the same direction. The DiD specification applied to $\log b^{(2)}$, the second-lowest bid in the iterative Pregão phase and the price-forming object in the IPV interpretation, delivers -0.0390 (SE 0.0206) in non-pharmaceuticals and -0.1307 (SE 0.0351) in pharmaceuticals. These estimates are built from raw bid units and do not pass through the reference-price normalization, so they corroborate the order-statistic channel directly—significantly in pharmaceuticals and marginally so in non-pharmaceuticals. Online Appendix OA-D.4 reports the full event study.

Winning prices also move. Table 2 reports the DiD on $\log p^{\text{final}}$ comparing Group 65 to never-treated product groups, with item and month fixed effects, auction-format and quantity controls, PBU controls, and item-clustered standard errors. The coefficient is written in a "DiD-in-reverse" convention: a negative coefficient on Group 65 interacted with the pre-period means that the open regime was cheaper, so the post-cutoff SME-only extension undoes that discount. With this sign convention, negative coefficients correspond to a post-cutoff price increase relative to the open-regime benchmark. Across windows, the log-price coefficient is between -0.108 and -0.148 with PBU controls; the 18-month benchmark is -0.113. This coefficient is used as a directional and timing check rather than as the source of the structural magnitude; the latter is built in Section 4 in p^{ref}-normalized units, and Section 6.6 examines reference-price evolution directly.

Table 2: Reduced-form price benchmark

Window around cutoff	6 months	12 months	18 months
$g65 \times Pre$ on $\log p^{\text{final}}$	-0.148	-0.108	-0.113
Standard error	(0.014)	(0.013)	(0.012)
Observations	201,588	418,593	625,598

Estimates use item and month fixed effects, controls for sealed-bid format and log quantity, PBU controls, and item-clustered standard errors. The coefficient is reported in the same sign convention: a negative pre-period coefficient is the open-regime price discount that disappears after the SME-only extension. The DiD is used as a sign, timing, and scale benchmark; the structural sections estimate the mechanism.

The reduced form establishes timing and direction. Magnitude is estimated structurally in Section 4, with the identification architecture set out in Section 2.4; OA-B reports the balance, placebo, and alternative-DiD diagnostics that discipline the reduced-form benchmark.

2.4. Identification architecture

The reduced-form design is not used as a stand-alone causal estimate of the legal reinterpretation. Its role is narrower: to verify that the institutional episode left the predicted footprint in timing, sign, and approximate scale. The structural decomposition and the static welfare comparison do the policy work.

The argument rests on three pieces of evidence that carry separate identifying weight. The difference-in-differences against 76 never-treated groups validates a real first-stage footprint around March 2018; the structural decomposition, conditional on observed entry and recovered willingness-to-supply primitives, identifies the price-forming mechanism; the static welfare comparison of V_0 against V_3, conditional on recovered primitives, maps those primitives into a static policy comparison. The three pieces rely on distinct assumptions and discipline different parts of the argument. A weakness in one piece does not mechanically overturn the others, but it changes which interpretation the evidence can support.

The reduced form requires that the legal episode be the dominant shock to Group 65 around March 2018 for the purpose of timing and sign; it is not asked to carry the structural magnitude. Heterogeneity-robust estimators attenuate and lose precision relative to TWFE: BJS imputation gives -0.0560 and Callaway–Sant'Anna gives -0.0165, the latter too imprecise to pin down the sign on its own. The reduced-form sign rests on TWFE and BJS and is corroborated directly by the raw second-lowest-bid DiD (Section 2); the structural magnitude is identified separately and does not lean on the DiD point estimate. The structural decomposition does not require cross-group balance because it conditions on within-period type-specific willingness-to-supply distributions, interpreted as cost distributions under the maintained model, and on equilibrium bidder counts; it requires the maintained IPV interpretation of Pregão drop-outs, disciplined in Section 6. The welfare comparison requires neither cross-group balance nor the DiD magnitude; it requires that the recovered primitives be interpretable for the static benchmark.

The diagnostics target six threats. Pre-treatment imbalance between Group 65 and the pooled controls limits the force of the reduced form (OA-B reports the balance table), but does not affect the structural exercise, which uses no cross-group comparison. Modern-DiD estimators attenuate and lose precision (BJS -0.0560; Callaway–Sant'Anna -0.0165, imprecise), consistent with the conservative reading assigned to the reduced form (OA-B). Reference-price evolution is examined directly in Section 6.6 and Online Appendix OA-D.4: a static DiD on $\log p^{\text{ref}}$ moves by -0.0273 in non-pharmaceuticals and -0.0325 in pharmaceuticals, below the corresponding movement in $\log p^{\text{final}}$, and an internal placebo confirms that the headline price effect survives where p^{ref} is stable. The IPV interpretation is the maintained restriction, disciplined in Section 6 and OA-C through censoring, common-scale, and cross-modality diagnostics. The empirical cutoff is the mass-take-up date rather than a single legal-document date; Section 2 documents the institutional sequence, and the phased-adoption sensitivity in OA-B excludes platform enablement months and retains the qualitative result. Strategic threshold manipulation through notice-splitting is addressed by a McCrary density test at the R$ 80,000 item-eligibility threshold (Section 6.7 and OA-D.4). Sections 6.3, 6.4, and 6.5 extend these checks to protected-pool composition, bidder coordination, and welfare-ranking sensitivity respectively, which refine rather than substitute for the six identification threats above.

The exercise does not identify a free-entry primitive, a structural model of SME selection, or dynamic capacity gains. Entry rates are observed equilibrium objects in the spirit of Athey et al. (2011), and the post-policy SME willingness-to-supply distribution is an observed active-bidder object, interpreted as a cost distribution under the maintained model, with strict invariance reported as a bounding exercise in Section 6.

3. Model and Identification

Section 2.4 laid out how the structural decomposition fits within the broader identification architecture. This section specifies the auction environment, the recovery of willingness-to-supply primitives, and the counterfactual price-formation objects. The model is deliberately spare. The empirical object is auction-scale willingness to supply revealed by exits; the structural object is the cost distribution obtained by interpreting those exits through the maintained IPV clock model. Observed entry then enters as an equilibrium object, and three counterfactual bidder pools—S_1, S_2, and S_3—isolate how the set-aside changes the price-forming order statistic.

3.1. Auction environment

Consider a single procurement auction for one indivisible contract. Each firm i has type $k_i \in \{\text{SME}, \neg\text{SME}\}$ and draws a private supply cost c_i from a type-specific distribution F_c^k. The two distributions are allowed to differ. The buyer runs an electronic Pregão, an English-reverse auction in which the price clock moves downward from the buyer's reference price and firms remain active until they are no longer willing to supply at the current price. The last active firm wins and supplies the contract at the clearing price.

Under the maintained IPV clock interpretation, a losing bidder's exit price is treated as its willingness to supply at the auction scale, following the English-auction logic in Milgrom and Weber (1982) and Haile and Tamer (2003). In the benchmark IPV clock model that dates back to Vickrey (1961), exit at cost is the relevant bidding rule; empirically, I use losing exits as willingness-to-supply observations and interpret them as costs under the maintained model. This is the key identifying advantage of the setting, but it is also the main structural restriction in the paper. In first-price procurement auctions the researcher observes strategically shaded bids and must recover costs by inverting a bid function under maintained value-paradigm restrictions (Guerre et al., 2000; Bajari and Hortaçsu, 2005). Here the price-forming statistic can be recovered from exits only if drop-outs are informative about willingness to supply rather than about collusive rotation, budget constraints, or other non-cost stopping rules.

The winner is censored because the winner does not drop out. Its cost is known only to be no higher than the second-lowest exit, or the corresponding final-price bound used in the data. The baseline estimator treats the winner's final price as a tight upper-bound observation; Section 6 reports both a losers-only estimator and a Turnbull NPMLE that treats the winner as left-censored at that bound. The expected transaction price in the clock interpretation is the expected second-lowest cost, $E[c_{(2)}]$, while the production cost of the allocation is the winner's cost, $E[c_{(1)}]$. This distinction matters later for welfare: $c_{(2)}$ determines the government's payment, while $c_{(1)}$ determines allocative efficiency.

3.2. Recovering willingness-to-supply primitives

The structural sample is organized by class, period, and bidder type. For each cell, loser drop-outs identify empirical distributions of normalized willingness to supply, which the maintained IPV clock model interprets as cost distributions. Athey and Haile (2002) provide conditions under which loser-exit data identify the underlying value distributions in ascending auctions, including the IPV-clock case used here, and Athey and Haile (2007) survey the broader nonparametric recovery toolkit. I first normalize exit prices by the buyer's reference price, so the object entering the simulation is measured in units of p_t^{ref}. This normalization absorbs the main cross-item scale differences, but it does not remove all auction-level heterogeneity. Two auctions in the same product class can differ in volume, delivery burden, regulatory documentation, or item-specific complexity in ways that move every bid in the auction together.

I therefore use a multiplicative auction-level heterogeneity correction in the spirit of

Krasnokutskaya (2011). Let

$$y_{it} = \log(b_{it}/p_t^{\text{ref}}) = a_t + e_{it},$$

where a_t is an auction-level scale component and e_{it} is the bidder-specific component. A method-of-moments variance decomposition estimates the between-auction and within-auction components, and a best-linear-predictor shrinkage removes $\hat{a}_t$ from each observation. The simulation uses $\exp(\hat{e}_{it})$ as the cleaned willingness-to-supply distribution, interpreted as a cost distribution under the maintained model. The correction removes common auction-level scale variation before comparing type-specific order statistics; it does not purge bidder-type selection, product-level composition, or post-policy changes in the active SME pool. In the current estimates, the resulting Pregão intraclass correlations range from 0.36 to 0.59 in the main cells, with values comparable to procurement settings where this correction is standard.

Section 6 disciplines this recovery step. It reports alternative treatments of winner censoring, a cross-modality check comparing Pregão drop-outs with a GPV recovery from Convite first-price bids (Guerre et al., 2000), dependence sensitivities, and bidder-pair coordination screens. These checks do not prove IPV. They ask the narrower question needed for this paper: whether the exclusion-dominant decomposition is mechanically produced by winner censoring, common auction-level scale heterogeneity, an obvious conflict between auction formats, or a post-policy increase in bidder clustering among the surviving SME pool. Online Appendix OA-C reports the corresponding quantile, KS, and censoring diagnostics.

3.3. Entry and the post-policy SME distribution

Entry is treated as observed equilibrium entry. The legal shock directly changes admissibility and observed participation; it does not identify fixed costs, bidder-specific participation rules, or a free-entry primitive. I therefore use the observed class-period type-specific participation margins to generate bidder counts in the simulations. In the baseline simulations, SME and non-SME bidder counts are drawn from independent Poisson processes whose means are set to the observed class-period type-specific average participation rates. The Poisson choice imposes mean-equals-variance for bidder arrivals at the cell level. Online Appendix OA-D.2 reports a robustness exercise that replaces the Poisson draws with the empirical class-period-type count distributions. The exclusion-dominant decomposition survives: net effects attenuate by roughly a quarter relative to the Poisson baseline (0.171 in non-pharmaceuticals and 0.223 in pharmaceuticals, against baseline 0.227 and 0.309), and exclusion shares move to 69.4 and 63.1 percent. The Poisson specification therefore slightly overstates the spread of bidder counts relative to the data, but the ranking of the two channels is unchanged. In S_1, both SMEs and non-SMEs enter according to their observed pre-policy participation rates. In S_2, non-SMEs are removed while the pre-policy SME participation rate is held fixed. In S_3, non-SMEs are removed and the observed post-policy SME participation rate is used.

This choice keeps the identifying burden where the data can bear it. A full entry model would require assumptions on fixed costs, information, bidder-specific participation decisions, and how the legal shock changes expected profits (Bajari et al., 2014). The legal shock identifies admissibility and observed participation, not a free-entry primitive. I therefore follow the reduced-form-entry spirit of Athey et al. (2011): observed bidder counts are equilibrium objects used in counterfactual price formation.

The same caution applies to the post-policy SME willingness-to-supply distribution, interpreted as the active-bidder cost distribution under the maintained model. The baseline simulation estimates $F_c^{\text{SME,Post}}$ directly from post-policy drop-outs and uses it as the distribution of active post-policy SME bidders. This object may reflect selection into the protected market, sourcing or technology changes within the SME segment, product-mix changes, or other composition changes in the active bidder pool around the cutoff. The baseline does not derive that difference from a formal selection model. The strict-invariance benchmark instead imposes $F_c^{\text{SME,Post}} = F_c^{\text{SME,Pre}}$. The comparison between these readings is a scope condition: the non-pharmaceutical policy ranking is stable, while the pharmaceutical ranking is not.

3.4. Counterfactual price-formation objects

The decomposition uses three simulated auction environments. These scenarios are not alternative legal regimes solved in equilibrium. They are counterfactual price-formation objects designed to isolate which bidder pool supplies the price-forming order statistic.

- S_1: open auction, pre-policy bidder pool. SMEs and non-SMEs enter according to their observed pre-policy Poisson participation rates, with pre-policy willingness-to-supply primitives.
- S_2: SME-only auction, pre-policy SME pool. Non-SMEs are removed, while the pre-policy SME participation rate and pre-policy SME willingness-to-supply primitives are held fixed.
- S_3: SME-only auction, observed post-policy SME pool. Non-SMEs are removed, and SMEs enter with the observed post-policy participation rate and the baseline post-policy willingness-to-supply primitives.

For each scenario, I draw bidder counts from the corresponding type-specific Poisson process, draw willingness-to-supply values from the recovered distributions (interpreted as costs under the maintained model), sort the draws, and record $c_{(1)}$ and $c_{(2)}$. The simulated price is $p_S = E_S[c_{(2)}]$, normalized by the reference price. The simulated production cost is $E_S[c_{(1)}]$.

The total set-aside price effect is

$$\Delta^{SA} = p_{S_3} - p_{S_1}.$$

It decomposes mechanically into two terms:

$$p_{S_3} - p_{S_1} = \underbrace{p_{S_2} - p_{S_1}}_{\Delta^{excl} \text{ lost competitive discipline}} + \underbrace{p_{S_3} - p_{S_2}}_{\Delta^{pool} \text{ protected-pool offset}} . \tag{1}$$

The first term holds the SME pool fixed and removes non-SMEs. It captures the price-forming cost of bidder exclusion: the second-order statistic is now drawn from the protected pool rather than from the full pool. The second term replaces the fixed pre-policy SME pool with the observed post-policy SME pool. It captures the extent to which additional protected participation and changes in active SME composition offset the exclusion shock. This term can be negative, and in the data it is: the realized protected pool partially restores competition, but not enough to undo the price increase. I refer to it as the protected-pool offset because the legal shock does not separately identify a pure entry primitive from selection into the active protected-pool composition.

The labels in equation (1) are order-statistic labels, not behavioral ones. In the maintained IPV clock model, bidders exit at willingness to supply, so the lost-discipline component is an admissibility effect rather than evidence of less aggressive SME bidding. This separates the channel studied here from first-price settings, where bidder preferences combine cost differences, order statistics, and strategic shading.

3.5. Scope of identification

The decomposition identifies counterfactual price-formation objects conditional on the maintained IPV clock interpretation, recovered willingness-to-supply primitives, and observed equilibrium entry. It does not identify a free-entry primitive, a structural model of SME selection, or dynamic capacity gains.

The model-specific restrictions are four. First, losing exits must be informative about willingness to supply rather than collusive rotation, budget constraints, or other non-cost stopping rules. Second, the post-policy SME willingness-to-supply distribution is an observed active-bidder object, not a primitive derived from a selection model. Third, the bidder-count process summarizes entry at the class-period-type level and abstracts from bidder-specific participation decisions. Fourth, bidder dependence and repeated-pair clustering remain first-order threats to the interpretation rather than cosmetic robustness concerns.

Section 6 disciplines these restrictions. It does not prove the maintained model. It asks whether the exclusion-dominant decomposition is overturned by the main threats: winner censoring, common auction-level scale heterogeneity, reference-price evolution, bidder dependence, or coordination among surviving firms.

4. Exclusion and the Price-Forming Pool

The set-aside changes the price-forming pool on two margins. It removes non-SMEs from eligible auctions and changes the protected SME pool through additional participation and composition. The first margin weakens the order statistic that disciplines price; the second pushes back. The evidence shows that the protected pool responds, but not enough to replace the competitive discipline supplied by excluded non-SMEs.

4.1. The protected pool responds

The first fact rules out the mechanical non-entry story. Protected firms do not fail to show up: SME participation rises sharply after the cutoff. In non-pharmaceutical Pregão auctions, the average number of SME bidders rises from 0.94 to 1.87. In pharmaceutical auctions, it rises from 0.55 to 1.22. Non-SME participation moves in the opposite direction, falling from 2.68 to 1.50 in non-pharma and from 2.61 to 1.66 in pharma.

That response matters because it works against the price increase. If the SME pool were held fixed at its pre-policy size and composition, the set-aside would look more expensive than the realized post-policy environment. The issue is not whether SMEs appear, but whether the resulting protected pool can replace the non-SME bidders that disciplined the open-auction price.

4.2. Exclusion dominates

Figure 1 puts the decomposition in one picture. The first bar, S_1, is the open pre-policy auction. The second, S_2, removes non-SMEs while holding the pre-policy SME pool fixed. The third, S_3, replaces that fixed pool with the observed post-policy SME pool. The jump from S_1 to S_2 is the lost competitive-discipline component. The decline from S_2 to S_3 is the protected-pool offset. The remaining gap between S_1 and S_3 is the realized price cost of the set-aside under the observed post-policy SME pool. These are counterfactual price-formation objects normalized by the reference price, not reduced-form treatment effects averaged over incomplete take-up.

Table 3 reports the same decomposition numerically. In non-pharmaceutical markets, moving from S_1 to S_2 raises the simulated price from 0.774 to 1.144, an increase of 0.371 of the reference price. Replacing the fixed pre-policy SME pool with the observed post-policy SME pool reduces the simulated price by -0.144. The net effect remains positive, at 0.227. In pharmaceutical markets, the exclusion component is larger in levels: 0.565 of the reference price, partly offset by -0.256, leaving a net effect of 0.309.

Because the protected-pool offset works in the opposite direction, I summarize dominance using both the exclusion-to-net ratio and the exclusion share of absolute component magnitudes. In non-pharmaceutical markets, the exclusion component is 164 percent of the net effect and 72.0 percent of the sum of absolute component magnitudes. Pharmaceutical procurement shows the same qualitative pattern at higher magnitudes (183 and 68.8 percent), under the scope conditions of Section 4.4. The dominant force

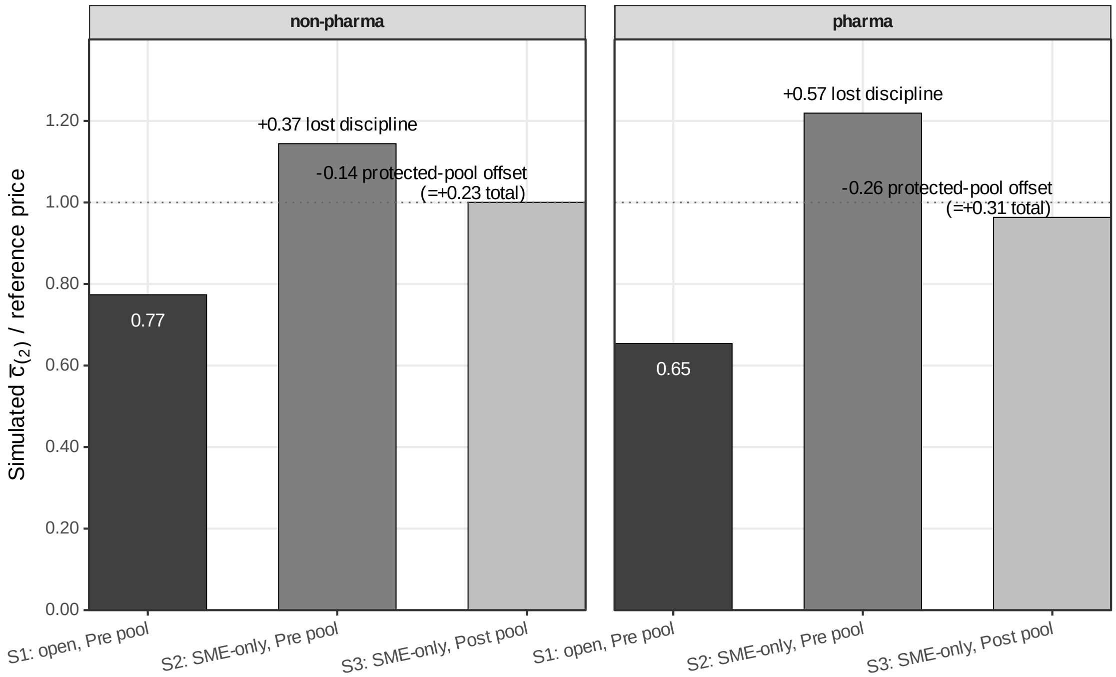

Figure 1: *Counterfactual price decomposition.* Bars report the simulated second-order statistic normalized by the buyer reference price. S_1 is the open pre-policy auction, S_2 removes non-SMEs while holding the pre-policy SME pool fixed, and S_3 allows the observed post-policy SME pool to replace the fixed pre-policy SME pool.

Table 3: Price decomposition: exclusion versus protected-pool offset

	Simulated price			Decomposition				
Class	S_1	S_2	S_3	$S_2 - S_1$	$S_3 - S_2$	$S_3 - S_1$	Excl./net	Abs. excl. share
Main result								
Non-pharma	0.774	1.144	1.000	+0.371	-0.144	+0.227	164%	72.0%
Boundary case (§4.4)								
Pharma	0.654	1.219	0.963	+0.565	-0.256	+0.309	183%	68.8%

S_1 is the open auction with the pre-policy bidder pool. S_2 is the SME-only auction holding the pre-policy SME pool fixed. S_3 is the SME-only auction using the observed post-policy SME pool. Prices are simulated $E[c_{(2)}]$ values normalized by the buyer reference price. "Excl./net" reports $(S_2 - S_1)/(S_3 - S_1)$. "Abs. excl. share" reports $|S_2 - S_1|/(|S_2 - S_1| + |S_3 - S_2|)$, where $S_3 - S_2$ is the protected-pool offset.

is therefore the loss of the non-SME price-forming pool, not a failure of protected firms to respond.

4.3. The protected-pool offset combines entry and composition

The term $S_3 - S_2$ should not be read as a pure entry parameter. It combines additional SME participation with changes in the active SME pool observed after the cutoff. In the baseline specification, $F_c^{\text{SME,Post}}$ is estimated from post-policy drop-outs and is therefore an active-bidder willingness-to-supply distribution, interpreted as a cost distribution under the maintained model. It may reflect selection into the protected market, sourcing or technology changes, product-mix changes, or other composition changes among active SME bidders. This is why I refer to $S_3 - S_2$ as a protected-pool offset rather than an entry effect.

Two checks discipline this interpretation. First, under the strict-invariance benchmark that forces the post-policy SME distribution to equal the pre-policy one, the exclusion component remains the larger component: 85 percent in non-pharma and 79 percent in pharma. Second, the Turnbull estimator gives the same qualitative message, with exclusion shares of 74.0 percent in non-pharma and 82.0 percent in pharma on the absolute-share normalization. Online Appendix OA-C and OA-D report the underlying censoring, primitive-stability, and mechanism tables.

The robust price-formation result is that bidder exclusion is the larger force in both classes. The policy implications, however, depend on whether the recovered primitives are stable enough for the welfare ranking to follow. Pharmaceutical procurement is the boundary case where they are not; Section 4.4 makes that scope condition explicit.

4.4. Pharmaceuticals as a boundary case

Pharmaceutical procurement is where the protected pool is thinnest, products are most heterogeneous, and post-policy composition turns over most. The logic of the decomposition predicts that the welfare ranking should be most fragile there, and it is.

Of post-policy pharmaceutical SME firms, 61.9 percent are new relative to the pre-period pool and account for 37.8 percent of post-policy SME bids, against 23.0 percent in non-pharmaceuticals. The pharmaceutical Pregão primitive-invariance diagnostic also fails after the heterogeneity correction (KS$= 0.1407$).

Under the strict-invariance benchmark in Section 5 and Online Appendix OA-E, the pharmaceutical welfare ranking reverses while the non-pharmaceutical ranking does not. That pattern is the scope condition itself. The price-formation result, that exclusion remains a large force, appears in both classes; the policy ranking, however, survives only where the recovered primitives are stable enough for the welfare comparison to discipline the protected pool. Pharmaceuticals are therefore reported as a boundary case; the headline result is non-pharmaceutical Group 65.

4.5. Interpretation

The mechanism is not that protected firms strategically bid less aggressively after the set-aside. In the maintained Pregão/IPV model, exits reveal willingness to supply, interpreted as cost under the maintained model. The mechanism is that the set-aside changes which order statistic forms the price. In the open auction, non-SMEs remain in the pool and discipline the second-lowest willingness-to-supply draw. In the SME-only auction, that discipline is removed. The post-policy protected pool pushes back against the loss, but it does not recreate the full-pool benchmark.

This price-formation result carries directly into welfare. A policy that supports SMEs by excluding non-SMEs buys redistribution by weakening the auction's price-forming pool. A policy that supports SMEs while leaving non-SMEs in the auction works through a different channel: it tilts allocation while preserving the bidders that discipline the price.

5. Static Welfare and Policy Design

I now use the recovered primitives for a static welfare accounting exercise and a static policy-design benchmark. The price decomposition maps naturally into welfare. The simulated second-order statistic, $c_{(2)}$, determines the government's payment; the winner's cost, $c_{(1)}$, determines allocative efficiency. A set-aside can raise payments without turning the full payment increase into deadweight loss: part of the extra payment is a transfer to the protected winner. The social cost is the real production inefficiency from assigning the contract to a higher-cost firm plus the marginal cost of financing the additional public outlay.

5.1. From price formation to welfare

Let p^V denote the expected procurement payment under policy V, and let $c_{(1)}^V$ denote the expected production cost of the winning supplier. Relative to the open benchmark S_1, the static per-auction welfare loss from the full SME-only set-aside V_0 is

$$\text{Loss}(V_0) = \underbrace{c_{(1)}^{V_0} - c_{(1)}^{S_1}}_{\text{DWL}_{\text{alloc}}} + \underbrace{\lambda \left(p^{V_0} - p^{S_1} \right)}_{\text{MCPF distortion}}. \tag{2}$$

The first term is real allocative waste: the contract is assigned to a supplier with a higher model-based production cost than under the open benchmark. The second term is the fiscal distortion from raising the additional public funds required to finance the higher expected payment. This MCPF accounting follows the public-economics convention going back to Ballard et al. (1985) and updated in modern welfare frameworks by Hendren and Sprung-Keyser (2020) and Finkelstein and Hendren (2020). The baseline uses $\lambda = 0.30$. The part of the payment increase not counted as allocative waste or fiscal distortion is an implicit transfer to the winning supplier. It enters social welfare as a benefit only to the extent that the planner assigns additional welfare weight to SME producer surplus.

5.2. Set-aside versus preference

Table 4 compares the full set-aside with a non-exclusionary static design benchmark. The comparison is not an implementation forecast. It holds entry and recovered willingness-to-supply primitives fixed to isolate the static value of preserving the non-SME price-forming pool. The full set-aside, denoted V_0, is the SME-only regime under the observed post-policy SME pool. The benchmark, denoted V_3, is a static open-auction scoring rule: SMEs receive a 10 percent price preference for winner selection, all firms remain eligible, and the government pays the actual winning bid. This benchmark should be read as a decomposition device, not as a forecast of a legal preference regime. It asks how much static welfare cost is avoided when SME support tilts allocation without deleting the non-SME bidders that discipline the price.

Table 4: Welfare cost of exclusion and the price-preference benchmark

	Full SME set-aside V_0				10% preference V_3	
Class	Δ_{gov}	$\text{DWL}_{\text{alloc}}$	MCPF	Loss / p_{S_1}	Δp	SME win gain
Non-pharma	0.247	0.148	0.074	28.9%	-0.004	4.3 pp
Pharma	0.298	0.207	0.089	44.8%	+0.002	1.4 pp

V_0 is the full SME-only set-aside under the observed post-policy SME pool. Δ_{gov}, $\text{DWL}_{\text{alloc}}$, MCPF, and Δp are reported in reference-price units. MCPF uses $\lambda = 0.30$. Loss / p_{S_1} is computed at full precision; the displayed Δ_{gov}, $\text{DWL}_{\text{alloc}}$, and MCPF components are rounded and need not reproduce the reported percentage exactly. V_3 is the static 10 percent preference benchmark described in the text. SME win gain is relative to the open benchmark S_1.

The full set-aside is costly on both welfare margins. In non-pharmaceuticals, it raises government payments by 0.247 of the reference price, of which 0.148 is allocative waste and 0.074 is the MCPF distortion. The total welfare loss is 28.9 percent of the open-regime price. In pharmaceuticals, the corresponding welfare loss is 44.8 percent. The lower endpoints of the bootstrap intervals remain economically large (20.5 percent in non-pharmaceuticals and 34.9 percent in pharmaceuticals).

The preference benchmark preserves the price-forming pool. Non-SMEs continue to participate and remain available to discipline the price-forming order statistic. Holding entry and recovered willingness-to-supply primitives fixed, the simulated price effect is essentially zero: -0.004 of the reference price in non-pharmaceuticals and +0.002 in pharmaceuticals. The preference is not a free version of the set-aside; it delivers less redistribution. The SME win-rate rises by 4.3/1.4 percentage points relative to the open benchmark rather than being forced to 100 percent. The comparison is therefore a frontier, not an unconditional ranking: set-asides buy more redistribution by removing rival bidders, while preferences buy less redistribution while preserving the price-forming pool.

The low static cost of the preference benchmark is not confined to the 10 percent case. In the preference grid, simulated welfare costs remain economically negligible up to 15 percent in non-pharmaceuticals and 25 percent in pharmaceuticals, and even a 30 percent preference costs only 4.99/1.92 percent of p_{S_1}. These are static comparisons, not implementation forecasts, but they show why exclusion is an extreme point on the policy frontier in thick standardized procurement. Online Appendix OA-E reports the full preference grid and the welfare-ranking sensitivity across the λ grid.

5.3. How much must the planner value SME surplus?

The set-aside is preferable to the 10 percent preference only if the planner places enough additional welfare weight on SME surplus. Let SMESurplus(V) denote expected SME producer surplus under policy V, computed from the simulated payment-cost wedge for SME winners. The indifference value is

$$w_\star^{\text{SME}} = \frac{\text{Loss}(V_0) - \text{Loss}(V_3)}{\text{SMESurplus}(V_0) - \text{SMESurplus}(V_3)}, \tag{3}$$

in the social-marginal-weight sense of Saez and Stantcheva (2016).

This is an indifference threshold, not an estimate of the planner's actual welfare weights. Within the static auction margin measured here, the implied threshold is 2.42

in non-pharmaceuticals and 2.61 in pharmaceuticals: to prefer exclusion to the 10 percent preference, the planner would need to value an extra real of SME surplus at more than twice a real of general surplus. Dynamic benefits of SME market access, learning, or capacity gains are additional terms on the benefit side and are not identified by the static auction exercise. Pharmaceutical sensitivity to the strict-invariance benchmark (0.7) is the same scope condition discussed in Section 4.4; the non-pharmaceutical ranking is stable.

5.4. Scope of the welfare ranking

The welfare ranking is conditional on the static auction primitives recovered above. These estimates do not imply that set-asides are never defensible. They imply that the defense must be explicit. Exclusion becomes more plausible when the protected pool is thin, when the planner wants protected market access rather than merely more marginal SME awards, when dynamic SME capacity gains are expected to be large, when goods are heterogeneous, or when the planner places a high social weight on SME surplus. Those conditions are not identified by the legal threshold alone. They are market-structure and welfare-weight conditions.

Non-pharmaceutical Group-65 procurement is the hard case for exclusion: the SME pool is thick, goods are standardized, and the preference benchmark preserves the price-forming pool while delivering SME wins at low static welfare cost. Pharmaceuticals are the boundary case: the SME pool is thinner, composition changes more under the policy, and the welfare ranking depends on how the post-policy SME willingness-to-supply distribution is interpreted. The design principle is therefore conditional, not universal. Exclusion is warranted only when the planner is willing to pay for the missing competitive discipline; otherwise, SME support should keep rival bidders in the auction.

6. Robustness and Threat Assessment

The checks below do not prove the maintained auction model. They ask the narrower question required by the paper: whether the exclusion-dominant decomposition is overturned by the main threats to its interpretation. I examine reduced-form timing and scale, winner censoring, auction-level heterogeneity, protected-pool composition, bidder dependence and coordination, the static policy benchmark, reference-price evolution, and threshold manipulation. No single check is decisive, and several leave real residual concerns. What the diagnostics together show is narrower: the main price-formation conclusion is not mechanically driven by any one of these threats.

The IPV-clock interpretation of Pregão drop-outs is load-bearing for the decomposition, so I evaluate it on four independent margins, reported in the relevant subsections below. (i) Winner censoring is varied across losers-only, all-bidders, and Turnbull NPMLE treatments; the exclusion share is stable at 74.0–82.0 percent across classes (Section 6.2). (ii) The Krasnokutskaya-style heterogeneity correction is paired with a Gaussian-copula dependence relaxation up to $\rho_c = 0.3$; the exclusion share drifts by less than 5 percentage points and the total effect by less than 10 percent across the grid (Section 6.2). (iii) The Convite first-price GPV recovery and the Pregão drop-out recovery line up in the load-bearing pharmaceutical non-SME pre-period cell after that correction (Section 6.2, Online Appendix OA-C). (iv) The Conley–Decarolis close-pair screen and the Bajari–Ye T1 ratio detect baseline within-class clustering but no post-policy intensification (Section 6.4), and the class-conditional persistent-pair statistic is at or below its within-class permutation null (Online Appendix OA-D). These four margins are not redundant: each addresses a different deviation from the IPV-clock benchmark. Their convergence is the strongest within-project discipline on the maintained interpretation; cross-jurisdictional replication remains the only path beyond it.

6.1. Reduced-form timing and scale

The reduced form is used for timing, sign, and approximate scale, not for the structural decomposition or welfare ranking. The 18-month benchmark in Table 2 is -0.113 in the paper's DiD-in-reverse convention, implying a 10–11 percent price movement. The 12- and 18-month estimates are close (-0.108 and -0.113); the 6-month window is somewhat larger in magnitude (-0.148). All three agree in sign and economic order of magnitude, which is the role the reduced form is asked to play. Excluding the BEC enablement months before the mass-take-up cutoff changes the coefficient from -0.109 to -0.087 with a standard error of 0.012, and placebo cutoffs before the reform are smaller in magnitude (-0.013, -0.030, and -0.034) than the real-cutoff estimates.

The conservative reading remains necessary. Group 65 is not balanced enough against the pooled controls to support a reduced-form-only design: 7 of 9 pre-period covariates exceed the 0.10 normalized-difference threshold, and 4 exceed 0.25. Modern DiD variants also attenuate the point estimate and reduce precision: the BJS imputation estimate (Borusyak et al., 2024) is -0.0560 and the Callaway–Sant'Anna estimate (Callaway and Sant'Anna, 2021) is -0.0165, imprecise enough that it does not by itself pin down the sign. The reduced-form sign is corroborated independently by the raw second-lowest-bid DiD (Section 2), which does not pass through the reference-price normalization. The attenuation pattern is consistent with the heterogeneity diagnostics of de Chaisemartin and D'Haultfoeuille (2020), Goodman-Bacon (2021), and Sun and Abraham (2021), all of which predict TWFE-versus-modern-DiD divergence when treatment timing or effects vary across units. The reduced form is therefore robust enough for timing and sign, but not strong enough to carry the structural magnitude or the welfare ranking. The paper uses it as external validation and estimates the mechanism with the auction decomposition. The structural magnitudes are not intended to match the reduced-form ATT: the DiD averages over incomplete take-up, product-time variation, and realized implementation, whereas S_1, S_2, and S_3 are counterfactual price-formation objects that isolate how the order statistic changes when the eligible bidder pool is altered. Online Appendix OA-B reports the balance table, placebos, and alternative estimators.

6.2. Willingness-to-supply recovery and auction primitives

The recovery concern is that the decomposition could be an artifact of how exits are converted into model-based cost primitives: winner censoring, common auction-level shocks, format-specific recovery, or bidder dependence. The diagnostics do not support that reading. Alternative treatments of winner censoring give similar net price effects: 0.275/0.259/0.246 in non-pharmaceuticals under losers-only/all-bidders/Turnbull inputs, and 0.347/0.308/0.357 in pharmaceuticals. The Turnbull exclusion share remains large (74.0 percent in non-pharmaceuticals and 82.0 percent in pharmaceuticals), so the result is not driven by treating the winner's final bid as an exact cost observation.

The other diagnostics point the same way. The auction-level heterogeneity correction removes common scale shocks before simulation, which asks whether common auction-level scale variation mechanically produces the exclusion component; it does not. A Convite first-price GPV recovery and the Pregão drop-out recovery line up in the key pharmaceutical non-SME pre-period cell after that correction, providing a cross-modality discipline rather than a proof of the model. A Gaussian-copula relaxation allows within-auction cost correlation up to $\rho_c = 0.3$; across that grid, the exclusion share moves by less than 5 percentage points and the total effect by less than 10 percent. Tightening or loosening the $c_\epsilon \leq 3$ filter and using 18-, 12-, or 6-month windows also leave the exclusion component as the larger component in both classes. These diagnostics do not show that the recovered primitives are literal costs. They show that the exclusion-dominant decomposition is not mechanically generated by winner censoring, common scale shocks, obvious format-specific recovery failures, moderate dependence, or the main filter/window choices. Online Appendix OA-C reports the willingness-to-supply quantiles and primitive-stability diagnostics behind these summaries.

6.3. Protected-pool composition

The main modeling threat is that the protected-pool offset combines participation and composition. The term $S_3 - S_2$ does not separately identify an entry primitive; it replaces the fixed pre-policy SME pool with the observed post-policy active SME pool. The baseline post-policy SME willingness-to-supply distribution is therefore an active-bidder object, interpreted as a cost distribution under the maintained model, rather than a primitive from a selection model.

The strict-invariance benchmark sets $F_c^{\text{SME,Post}} = F_c^{\text{SME,Pre}}$ while keeping observed post-policy SME entry counts. The total price effect remains positive: 0.29 in non-pharmaceuticals and 0.47 in pharmaceuticals, and the exclusion shares rise to 85 and 79 percent. This check does not show that composition is irrelevant. It shows that composition is not what makes the exclusion component dominate the price decomposition. Composition matters most for the pharmaceutical welfare ranking, which Section 4.4 treats explicitly as a boundary case.

6.4. Coordination

Coordination is the conduct threat that would most directly contaminate the drop-out interpretation. If repeated bidder pairs coordinate exits, drop-out prices may reflect conduct rather than willingness to supply (Kawai and Nakabayashi, 2022). The screens detect baseline clustering; the identifying question for this paper is narrower: does clustering intensify after non-SMEs are removed?

It does not. The Conley-style close-pair screen (Conley and Decarolis, 2016) is stable in non-pharmaceuticals: realized shares are 16.9 percent before and 16.8 percent after the cutoff, compared with null means of 10.6 and 10.2. In pharmaceuticals, the realized share falls from 27.6 to 24.4 percent. Bajari–Ye-style ratios (Bajari and Ye, 2003) show the same directional pattern: the non-pharmaceutical ratio falls from 2.63 to 1.83, and the pharmaceutical ratio falls from 1.29 to 1.11.

These screens weaken the differential-coordination story, not the broader possibility of noncompetitive conduct. Residual baseline clustering remains a limitation. The diagnostics answer the narrower threat to the decomposition: the main post-policy price increase is not obviously driven by a new coordination shock among the surviving SME pool. Online Appendix OA-D reports the coordination-screen table.

6.5. Static policy benchmark

The price-preference comparison is a static design benchmark, not an implementation forecast. The robustness question is whether the welfare ranking that favors preserving the price-forming pool in standardized non-pharmaceutical markets survives changes in λ and in the treatment of the post-policy SME primitives. The non-pharmaceutical static welfare ranking remains $V_3 \succ V_0$ over $\lambda \in [0.15, 0.45]$ under both the main and strict-invariance specifications. In pharmaceuticals, by contrast, the ranking remains conditional on whether the post-policy SME pool is modeled directly or held invariant.[1] Online Appendix OA-E reports the preference grid and welfare-weight sensitivities.

6.6. Reference-price evolution

Because the structural objects are normalized by buyer reference prices, endogenous reference-price evolution at the cutoff could contaminate reduced-form interpretation and structural levels. A direct TWFE DiD on $\log p^{\text{ref}}$ moves by only -0.0273 (SE 0.0105) in non-pharmaceuticals and -0.0325 (SE 0.0124) in pharmaceuticals—small relative to the corresponding DiD on $\log p^{\text{final}}$ in each class (-0.0741 and -0.1444). The residual price effect after subtracting the p^{ref} component is about 4.7 percent in non-pharma and

[1] The result is not an artifact of exact Vickrey equivalence: a Maskin–Riley first-price upper-bound adjustment (Maskin and Riley, 2000) is at most 5 percent of the reference price, while the simulated V_0–V_3 price gap is 22–26 percent.

11.2 percent in pharma. Modern estimators (BJS imputation, Callaway–Sant'Anna) are consistent in sign; Online Appendix OA-D.4 reports the full estimator comparison and the underlying event study.

The internal placebo by quartile of $|\Delta \log p^{\text{ref}}|$ tightens the argument. Items in the lowest quartile, where the reference price was effectively stable across the cutoff, still show a price effect of -0.0438 in non-pharmaceuticals and -0.0923 in pharmaceuticals—about two-thirds of the class-specific DiD coefficient. The headline price movement is therefore not concentrated only among items whose reference prices moved.

A uniform p^{ref} shift would not mechanically generate the decomposition: the structural objects are constructed in p^{ref}-normalized units, and the willingness-to-supply recovery is unaffected by a common scale change. The relevant concern is heterogeneous reference-price movement correlated with the reform; the quartile placebo addresses that channel directly. Threshold manipulation, the canonical buyer-side channel through which reference-price gaming would operate, is examined separately in Section 6.7.

6.7. Threshold manipulation

The canonical buyer-side manipulation threat is threshold gaming: buyers could split or restructure purchases around the R$ 80,000 item-eligibility threshold. A McCrary density test at the threshold shows no detectable bunching of items (post: $T = 1.06$, $p = 0.29$; pre: $T = -1.43$, $p = 0.15$). Group 65 items also sit well below the R$ 80,000 item-eligibility threshold in mass (Online Appendix OA-B), so systematic manipulation around it should be visible if present.

The density evidence does not indicate systematic splitting or restructuring around the threshold. This does not rule out all buyer-side adaptation, but it rules against the canonical threshold-gaming channel of the type Coviello et al. (2022) document in procurement settings where manipulation is empirically active. The result is consistent with the reference-price evidence in Section 6.6: the mild p^{ref} movement detected there is not aligned with threshold gaming.

Table 5 summarizes the threat assessment. The diagnostics discipline the maintained interpretation; they do not prove the auction model.

Table 5: Threat assessment: diagnostics, findings, and residual limitations

Threat	Diagnostic	Finding	Remaining limitation
Weak reduced form	Balance checks, placebo cutoffs, modern DiD estimators	Timing and sign supported; magnitude attenuates under CS/BJS	Not a stand-alone causal estimate of the legal reinterpretation
Winner censoring	Losers-only, all-bidders, Turnbull NPMLE	Exclusion share remains large under each treatment	Winner costs are not directly observed
Bidder-count process	Poisson baseline; empirical class-period-type count distributions (OA-D.2)	Empirical-count robustness: exclusion shares 69.4/63.1 (NP/PH); ranking unchanged	No full entry model; counts are equilibrium objects, not primitives
Protected-pool composition	Strict-invariance benchmark	Exclusion still dominates the price decomposition	Pharmaceutical welfare ranking remains sensitive to composition
Coordination and conduct	Conley-style close-pair screens; Bajari–Ye ratios	No post-policy intensification of clustering	Baseline clustering remains; broader noncompetitive conduct not ruled out
Reference-price evolution	DiD on $\log p^{\text{ref}}$; stable-reference quartile placebo	Price effect survives where p^{ref} is stable	Heterogeneous p^{ref} movement cannot be fully ruled out
Threshold manipulation	McCrary density test at R$ 80,000 threshold	No detectable bunching of items at the threshold	Does not rule out all buyer-side adaptation
Static policy benchmark	λ grid; strict invariance; preference grid	Non-pharmaceutical static welfare ranking is stable	Not an implementation forecast; pharmaceutical ranking conditional

Each row maps one identification or sensitivity threat to the diagnostic that disciplines it. The findings column reports the direction of the diagnostic, not a proof of the maintained auction model or the welfare ranking. Cross-references are Section 6.1 (reduced-form), 6.2 (winner censoring), 6.3 (protected-pool composition), 6.4 (coordination), 6.5 (policy benchmark), 6.6 (reference-price evolution), and 6.7 (threshold manipulation).

Taken together, the checks do not prove the maintained auction model. They show that the exclusion-dominant decomposition is not overturned by the main mechanical threats: reduced-form timing, winner censoring, common scale shocks, protected-pool composition, bidder dependence, reference-price evolution, or threshold manipulation. The remaining interpretation is the one the decomposition emphasizes: when the excluded bidders are the ones disciplining the price-forming order statistic, set-asides are costly because they replace competition with eligibility.

7. Conclusion

This paper shifts the evaluation of SME set-asides from access alone to price formation. A set-aside does not merely direct contracts toward protected firms; it changes the bidder pool from which the auction price is formed. When the excluded non-SMEs are the bidders that discipline the relevant order statistic, SME access is purchased by weakening the competitive mechanism itself.

The paper does not estimate a full reduced-form treatment effect of the legal reinterpretation, nor does it solve a dynamic entry model of SME development. It identifies a narrower object: how the price-forming order statistic changes when an exclusionary rule removes non-SMEs and the protected SME pool responds. Within that static auction margin, exclusion is expensive in standardized non-pharmaceutical procurement. A planner can still prefer set-asides, but the justification must come from benefits outside the measured static margin: SME surplus weights, market-access objectives, or dynamic capacity gains.

The evidence from São Paulo's procurement reform shows that the protected pool responds, but incompletely. SME participation rises after the expansion of SME-only tendering, yet the post-policy protected pool does not recreate the price discipline supplied by excluded non-SMEs. In standardized non-pharmaceutical procurement, this leaves a large static cost: the full set-aside generates a welfare loss of 28.9 percent of the open-regime price at $\lambda = 0.30$. Pharmaceutical procurement exhibits the boundary of the exercise. There, the protected pool is thinner, composition changes more, and the welfare ranking becomes model-sensitive.

The policy implication is that the instrument matters. A full set-aside is a high-redistribution instrument because it forces eligibility by removing rival bidders. A price preference delivers less redistribution, but it preserves the non-SME bidders that discipline the price-forming order statistic. The relevant comparison runs not between SME support and no support, but between exclusionary redistribution and support that keeps the price-forming bidder pool inside the auction. The evidence does not imply that an implemented preference regime would reproduce the static benchmark exactly; it shows that exclusion is the costly margin, and that preserving the price-forming pool is valuable before any behavioral response is modeled.

In thick standardized markets, where a preference rule can tilt allocation without deleting rival bidders, exclusion is an expensive way to redistribute. Support for SMEs need not remove the bidders that discipline the price.

Acknowledgements

I thank Paulo Furquim de Azevedo for advising and continued discussion of this project, and seminar audiences at INSPER for comments on earlier drafts. All remaining errors are my own.

Declaration of competing interests

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The structural sample is derived from BEC administrative microdata accessed through a research agreement with the Secretaria da Fazenda e Planejamento do Estado de São Paulo. The raw administrative records are not publicly redistributable under the agreement. Replication materials with code, generated tables, figures, and non-confidential derived outputs will be made available in a public repository, subject to the data-use restrictions.

Declaration of generative AI in the writing process

During the preparation of this manuscript, the author used Anthropic Claude to assist with copy-editing and prose revision. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

References

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

Athey, S., Haile, P.A., 2002. Identification of standard auction models. Econometrica 70, 2107–2140.

Athey, S., Haile, P.A., 2007. Nonparametric approaches to auctions, in: Heckman, J.J., Leamer, E.E. (Eds.), Handbook of Econometrics. Elsevier. volume 6A. chapter 60, pp. 3847–3965.

Athey, S., Levin, J., Seira, E., 2011. Comparing open and sealed bid auctions: Evidence from timber auctions. Quarterly Journal of Economics 126, 207–257.

Bajari, P., Hortaçsu, A., 2005. Are structural estimates of auction models reasonable? evidence from experimental data. Journal of Political Economy 113, 703–741.

Bajari, P., Houghton, S., Tadelis, S., 2014. Bidding for incomplete contracts: An empirical analysis of adaptation costs. American Economic Review 104, 1288–1319.

Bajari, P., Ye, L., 2003. Deciding between competition and collusion. Review of Economics and Statistics 85, 971–989.

Ballard, C.L., Shoven, J.B., Whalley, J., 1985. General equilibrium computations of the marginal welfare costs of taxes in the united states. American Economic Review 75, 128–138.

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

Borusyak, K., Jaravel, X., Spiess, J., 2024. Revisiting event-study designs: Robust and efficient estimation. Review of Economic Studies 91, 3253–3285.

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

Callaway, B., Sant'Anna, P.H.C., 2021. Difference-in-differences with multiple time periods. Journal of Econometrics 225, 200–230.

Conley, T.G., Decarolis, F., 2016. Detecting bidders groups in collusive auctions. American Economic Journal: Microeconomics 8, 1–38.

Coviello, D., Guglielmo, A., Lotti, C., Spagnolo, G., 2022. Procurement with manipulation. CEPR Discussion Paper 17063. Centre for Economic Policy Research. Revise and resubmit, Journal of Public Economics.

Coviello, D., Mariniello, M., 2014. Publicity requirements in public procurement: Evidence from a regression discontinuity design. Journal of Public Economics 109, 76–100.

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

Decarolis, F., Fisman, R., Pinotti, P., Vannutelli, S., 2025. Rules, discretion, and corruption in procurement: Evidence from italian government contracting. Journal of Political Economy: Microeconomics 3, 213–254.

Ferraz, C., Finan, F., Szerman, D., 2016. Procuring firm growth: The effects of government purchases on firm dynamics. Working paper.

Finkelstein, A., Hendren, N., 2020. Welfare analysis meets causal inference. Journal of Economic Perspectives 34, 146–167.

Goodman-Bacon, A., 2021. Difference-in-differences with variation in treatment timing. Journal of Econometrics 225, 254–277.

Guerre, E., Perrigne, I., Vuong, Q., 2000. Optimal nonparametric estimation of first-price auctions. Econometrica 68, 525–574.

Haile, P.A., Tamer, E., 2003. Inference with an incomplete model of English auctions. Journal of Political Economy 111, 1–51.

Hendren, N., Sprung-Keyser, B., 2020. A unified welfare analysis of government policies. The Quarterly Journal of Economics 135, 1209–1318.

Kawai, K., Nakabayashi, J., 2022. Detecting large-scale collusion in procurement auctions. Journal of Political Economy 130, 1364–1411.

Krasnokutskaya, E., 2011. Identification and estimation of auction models with unobserved heterogeneity. Review of Economic Studies 78, 293–327.

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

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

Maskin, E., Riley, J., 2000. Asymmetric auctions. Review of Economic Studies 67, 413–438.

Milgrom, P.R., Weber, R.J., 1982. A theory of auctions and competitive bidding. Econometrica 50, 1089–1122.

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

Saez, E., Stantcheva, S., 2016. Generalized social marginal welfare weights for optimal tax theory. American Economic Review 106, 24–45.

Sun, L., Abraham, S., 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics 225, 175–199.

Vickrey, W., 1961. Counterspeculation, auctions, and competitive sealed tenders. Journal of Finance 16, 8–37.