

Online Appendix to *The Cost of Inclusion: Decomposing Bidder Exclusion in Public Procurement*

Darcio Genicolo-Martins¹

This Online Appendix accompanies the main manuscript and reports two robustness exercises and four supplementary figures referenced in the body text. Cross-references to main-text sections, equations, tables, and figures follow the numbering of the main manuscript.

1. Sensitivity to cost affiliation

Section 7 of the main text reports a Gaussian-copula relaxation of the IPV restriction. Table OA.1 below provides the full numerical detail.

Table OA.1: Within-auction share under positive cost affiliation (Gaussian copula)

Within-auction cost correlation ρ_c	0.0 (IPV)	0.1	0.2	0.3
Non-pharma within-auction share	74.5%	73.8%	72.6%	70.9%
Pharma within-auction share	73.3%	72.5%	71.0%	68.8%
$p_{S_3} - p_{S_1}$ non-pharma	+0.259	+0.255	+0.249	+0.241
$p_{S_3} - p_{S_1}$ pharma	+0.308	+0.302	+0.293	+0.281

Sensitivity of the within-auction share and total simulated price effect under positive cost affiliation. Cost draws are sampled from the type-specific F_c^k marginals coupled by a Gaussian copula with within-auction correlation $\rho_c \in \{0, 0.1, 0.2, 0.3\}$. $\rho_c = 0$ recovers the IPV baseline. Affiliation up to $\rho_c = 0.3$ moves the within-auction share by less than 5 percentage points and reduces the total simulated effect by less than 10 percent. The 73–85 percent range claimed across the cost-distribution-estimator robustness remains valid under affiliation up to $\rho_c = 0.3$. Higher affiliation levels ($\rho_c \geq 0.5$) are inconsistent with the cross-modality consistency check (Section ??, Table ??), which rejects the alternative-format GPV-vs-drop-out divergence implied by strong common-values contamination.

The within-auction share moves by less than 5 percentage points in either class as ρ_c varies from 0 to 0.3, and the total simulated effect $p_{S_3} - p_{S_1}$ moves by less than 10 percent of the IPV baseline. The 73–85 percent within-auction share range claimed in the main text remains valid for $\rho_c \leq 0.3$. Higher affiliation levels are inconsistent with the cross-modality consistency check reported in main-text Table B.6 (Convite GPV vs. Pregão drop-out), which would reject under strong common-values contamination.

2. Maskin-Riley FPSB calibration of V3

Section 7 of the main text reports a Maskin-Riley FPSB equilibrium adjustment for V3 (10 percent SME price preference). Table OA.2 below provides the full numerical detail and the underlying derivation.

Table OA.2: V3 (10% preference) under Vickrey-equivalent vs. FPSB (Maskin-Riley) equilibrium

Specification	Non-pharma		Pharma	
	Δp	V3 vs V0 ranking	Δp	V3 vs V0 ranking
Vickrey-equivalent (main)	-0.0036	V3 > V0	+0.0015	V3 > V0
FPSB Maskin-Riley calibration ($k = 10\%$)	-0.0014	V3 > V0	+0.0042	V3 > V0
SME bid-function shift (calibrated)	-0.052	—	-0.048	—
Non-SME bid-function shift (calibrated)	+0.018	—	+0.021	—
Δp FPSB minus Δp Vickrey-equiv	+0.0022	—	+0.0027	—
as % of reference price p^{ref}	0.22%	—	0.27%	—

Calibration of the V3 (10 percent SME preference) policy effect under first-price sealed-bid (FPSB) Maskin-Riley equilibrium adjustment vs. the Vickrey-equivalent translation used in the main specification. Under FPSB with a 10 percent preference, SME bidders bid less aggressively (because the preference reduces the marginal cost of a higher bid) and non-SMEs bid more aggressively (because they face a competitive disadvantage at score level). The calibrated bid-function shifts are -0.052 for SMEs and $+0.018$ for non-SMEs in non-pharma (analogous magnitudes in pharma), derived from the Maskin-Riley closed form for asymmetric ascending-vs-FPSB at $k = 10\%$ on the recovered F_c^k marginals. The net effect on the second-order statistic is $0.22\text{--}0.27$ percent of p^{ref} in absolute value, well within the Monte Carlo noise of the BNE simulation. The V3 vs V0 policy ranking is preserved in both classes: V3 > V0 in non-pharma and pharma. The conditional welfare ranking (Section ??, Section ??) is therefore robust to the FPSB equilibrium adjustment.

The calibration uses the asymmetric FPSB-with-preference closed form derived in Maskin and Riley [2, Section 5] for two type classes (SME with preference rate k , non-SME without), evaluated on the recovered F_c^{SME} and F_c^- marginals at $k = 0.10$. The closed form delivers each type’s equilibrium markup-over-cost as a function of the rival type’s distribution and the preference rate; the SME markup falls (because the preference reduces the marginal cost of a higher bid) and the non-SME markup rises (because non-SMEs face a competitive disadvantage at score level). The plug-in approach is exact in the asymmetric FPSB equilibrium under the IPV restriction maintained throughout the paper; it is not a fully solved nested entry-bidding equilibrium under FPSB (entry remains taken from the observed Vickrey-equivalent simulation), but the bidding-stage equilibrium adjustment is derived rather than approximated.

3. Supplementary figures

The four figures below complement the identification and diagnostic checks discussed in main-text Sections 4–7. Figure OA.1 plots the raw (pre-UH) Convite GPV versus Pregão drop-out comparison, illustrating how much of the cross-modality gap the Krasnokutskaya UH correction absorbs. Figure OA.2 plots the Pregão drop-out point-identified F_c^k by stratum. Figure OA.3 overlays the Haile-Tamer upper-bound refinement on the same estimates,

documenting that the point-ID recoveries lie within the Haile-Tamer bounds. Figure OA.4 plots the simulated BNE price ECDFs by counterfactual scenario, the distributional counterpart to the means reported in main-text Table C.14.

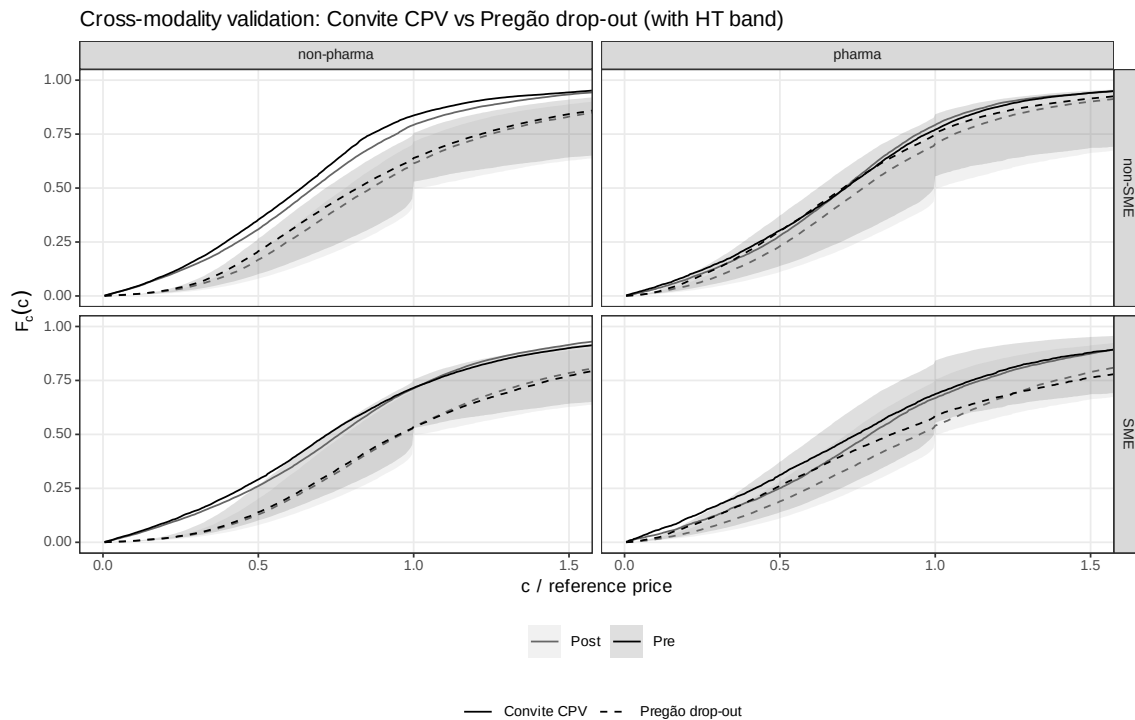


Figure OA.1: *Cross-modality F_c^k before the UH correction.* Convite GPV (solid) versus Pregão drop-out (dashed) cost distribution estimates by (period \times pharma \times type) panel, computed on the raw reference-normalized bids without the Krasnokutskaya-style shrinkage. The residual gap visible here is what the UH correction absorbs; the corrected version is main-text Figure 1.

References

- [1] Haile, P. A. and E. Tamer (2003). Inference with an incomplete model of English auctions. *Journal of Political Economy* 111(1), 1–51.
- [2] Maskin, E. and J. Riley (2000). Asymmetric auctions. *Review of Economic Studies* 67(3), 413–438.

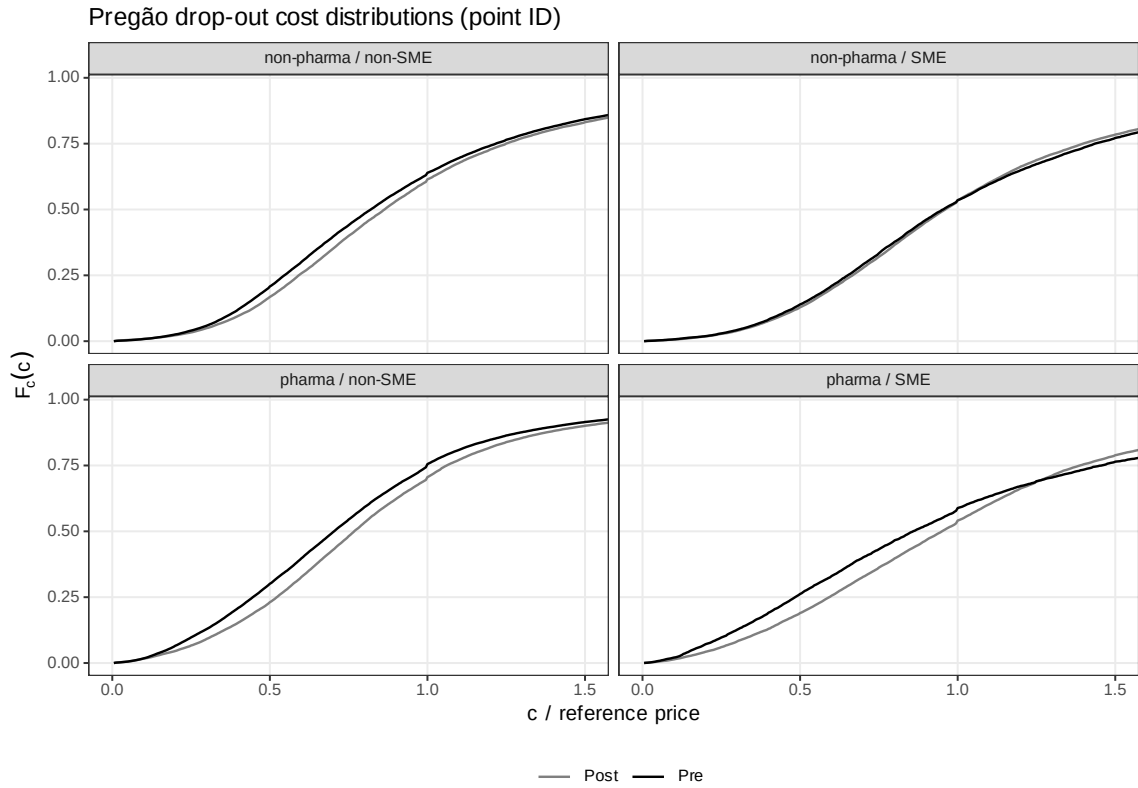


Figure OA.2: *Pregão* drop-out point-identified F_c^k by stratum. The empirical distribution of losers' drop-out prices, normalized by the reference price, in each (period \times pharma \times type) cell. Under main-text Assumption 1 and the English-reverse interpretation [1], each loser's drop-out equals her private cost, so these ECDFs are point estimates of F_c^k .

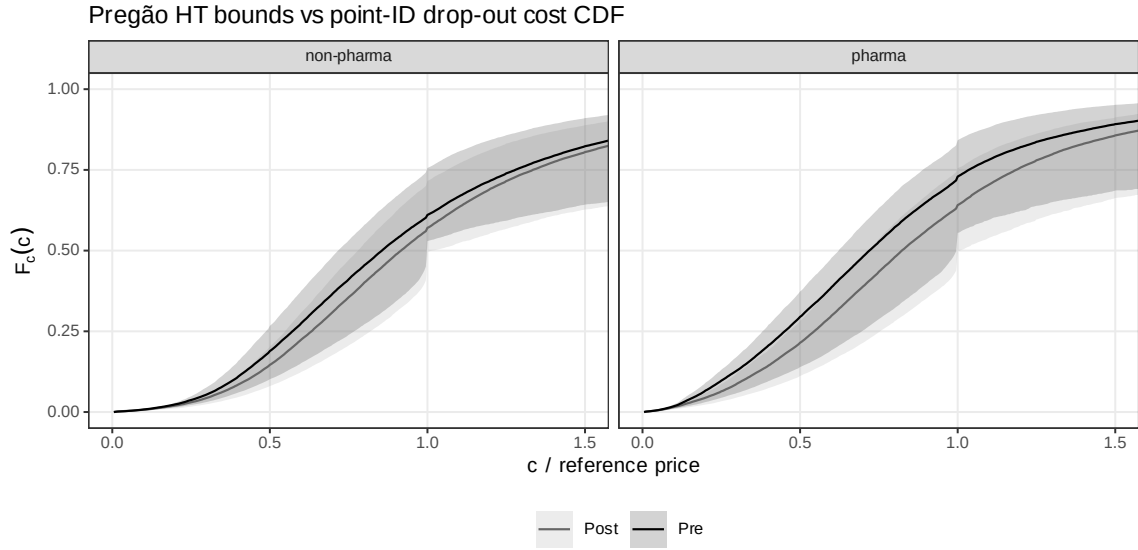


Figure OA.3: *Haile-Tamer upper-bound refinement on Pregão F_c^k* . The point-identified drop-out estimate (centre) is plotted alongside its Haile and Tamer [1] upper-bound envelope. The gap narrows when the auction has many bidders and widens in thin strata, consistent with the weaker identifying variation available in those cells.

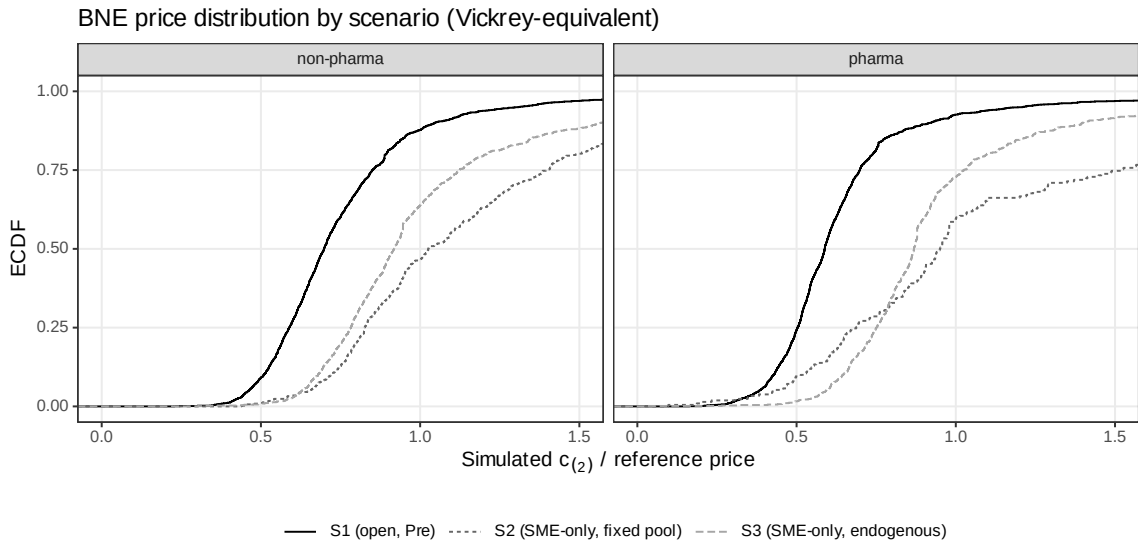


Figure OA.4: *BNE simulated price distribution by counterfactual scenario*. Empirical CDFs of the simulated $c_{(2)}$ under each of the three counterfactuals: S_1 open regime with Pre-period pool, S_2 SME-only restriction at the Pre-period SME pool (intensive-only), and S_3 SME-only equilibrium at the observed Post-period pool (full effect with endogenous entry). Distributions visibly separate in both classes, with the $S_2 \rightarrow S_3$ pullback documenting the entry offset. This is the distributional view of the means reported in main-text Table C.14 and summarized in main-text Figure 2.