Bitter Pills: Judicial Enforcement and the Cost of Public Procurement in Brazil*

Darcio Genicolo-Martins[a,*], Paulo Furquim de Azevedo[a]

[a] *Insper Institute of Education and Research, Rua Quatá 300, São Paulo, SP 04546-042, Brazil*

Abstract

Judicial enforcement is essential for the rule of law, yet the *design* of enforcement can generate economic inefficiency. We provide the first quantitative decomposition of the procurement cost of judicial enforcement, using bid-level data on court-mandated pharmaceutical procurement in São Paulo, Brazil (2009–2019). With item, time, and buyer fixed effects, court orders raise prices by 5.4% and reduce bidder participation by -5.4%. To isolate the role of personal sanctions specifically, we exploit an institutional feature unique to São Paulo since 2009: a parallel administrative-request channel that shares all planning constraints of court-mandated procurement but carries no penalty for officials when delivery fails. The "under the gun" comparison between litigated and administrative urgent purchases of the same

*Short paper. This version: May 2026. We thank [acknowledgments to be added]. All errors are our own.

*Corresponding author.

Email addresses: darciogm1@insper.edu.br (Darcio Genicolo-Martins), paulofa1@insper.edu.br (Paulo Furquim de Azevedo)

item yields a 23–30% cost gap. Decomposing this umbrella into its underlying channels, we find that demand fragmentation—smaller orders that forgo bulk discounts—accounts for essentially all of the under-the-gun gap and roughly half of the urgency premium across the full sample; reduced bidder participation, even at identical order size, accounts for most of the rest; and a within-firm markup under urgency is not what we find. The findings speak to a general asymmetry in accountability design: when non-completion is punished more severely than overpayment, compliance crowds out cost-efficiency.
Keywords: public procurement, health litigation, enforcement costs, judicial mandates, passive waste

JEL: D44, D73, H51, H57, I18, K41

1. Introduction

A judge orders the government to buy a cancer drug within 1–10 days. The procurement officer has no time to aggregate demand across other purchases, no time to research market prices, and faces personal fines if delivery fails. The drug arrives, but the cost margin between this transaction and the same agency's ordinary purchase of the same drug runs from a few percent at the typical urgent purchase to as much as a third more for purchases that are court-mandated rather than initiated through an internal administrative channel.

This paper provides the first quantitative decomposition of that cost margin. Courts and regulatory agencies are alternative instruments of social control (Glaeser and Shleifer, 2003), but court intervention in procurement compresses planning timelines and distorts incentives in ways that raise prices and reduce competition (Coviello et al., 2018b). In health care, the stakes are high: judicial mandates directly determine how—and at what price—governments acquire essential medicines (Wang, 2015; Ferraz, 2009). Brazil's 1988 Constitution enshrines health as a "right of all and a duty of the state," and courts have granted virtually every medication request—96,000 first-instance claims nationally in 2017, a 913% increase in São Paulo since 2008 (CNJ/INSPER, 2019).

We use bid-level data covering all pharmaceutical procurement by the São Paulo State Department of Health (SES/SP) from 2009 through 2019: over 479,330 purchase-offer-item observations conducted through the state's electronic procurement platform (BEC). Court orders and administrative requests partition purchases into *ordinary* (standard timeline) and *urgent* (compressed timeline). Within urgent purchases, a further institutional distinction proves useful: *administrative* requests, instituted by SES/SP in 2009, carry no penalty for procurement failure, while *litigated* purchases expose officials to fines, personal liability, and fund seizure. Comparing these two types—which share planning constraints but differ in penalty exposure—is what we call the "under the gun" (UTG) comparison. *Our identification is descriptive throughout. We decompose a within-item gap between two procurement regimes that share planning constraints but differ in penalty exposure; we do not estimate a counterfactual price under hypothetical sanction* *removal. That object is unidentified in our setting because the administrative subsample is itself selected by a scientific committee on cost-effectiveness criteria.* What we provide is a quantitative decomposition of where the cost margin between regimes comes from.

Across the full sample, urgent purchases carry negotiated prices 5.4% higher, reference prices 2.7% higher, and attract -5.4% fewer bidders—yet succeed *more often* (2.1 pp), consistent with officials accepting worse terms to avoid failure. The UTG comparison is sharper: litigated purchases cost 23–30% more than administrative ones in the same item under the same fixed effects. We then decompose this UTG umbrella into three distinct channels. *Demand fragmentation* (C1) does most of the work: orders shrink under judicial pressure, and the bulk-discount channel mechanically accounts for essentially all of the UTG gap and roughly half of the urgency premium across the full sample. *Reduced participation* (C2) explains the residual urgency premium across the full sample: even at identical order size, urgent tenders attract about -4.2% fewer bidders, leaving thinner markets and worse negotiated prices. *Supplier composition shift* (C3) explains the rest of the across-sample premium: the supplier mix differs between urgent and ordinary regimes; within firm-buyer-item triples observed in both regimes, the within-firm price under urgency is, if anything, slightly lower, ruling out a within-firm markup story.

Three contributions follow. First, we provide the first quantitative decomposition of the procurement cost of judicial enforcement, contributing to a large literature on government spending efficiency and procurement design (Bandiera et al., 2009; Best et al., 2023; Decarolis et al., 2020; Lewis-Faupel et al., 2016; Szücs, 2024). The decomposition relies on a methodological discipline that explicitly avoids post-treatment-bias: we read controls on quantity and supplier identity as channel-decomposing rather than as direct-effect estimators (Acharya et al., 2016). Second, the UTG comparison—exploiting within-urgent variation in penalty exposure—offers direct evidence that accountability mechanisms can distort bureaucratic decisions, echoing a long-standing theoretical prediction (Prendergast, 2007) and complementing recent field evidence on bureaucratic effort, autonomy, and management (Rasul and Rogger, 2018; Williams, 2021; Bandiera et al., 2021). Third, in contrast to Coviello et al. (2018b), where slow courts *extend* procurement timelines, we study injunctions that *compress* them, adding to the evidence on ex ante design in auction performance (Coviello et al., 2018a; Baltrunaite et al., 2021) and on competition design under contractual incompleteness (Carril et al., 2026). We complement the legal and public-health literatures on the judicialization of health in Latin America (Wang, 2015; Ferraz, 2009; Biehl et al., 2009), which have focused on legal and clinical dimensions with limited attention to procurement costs.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 presents results. Section 6 concludes.

2. Institutional Background

2.1. Health as a Constitutional Right and the Rise of Litigation

Brazil's 1988 Constitution created the Unified Health System (SUS), which maintains a formulary of approved medications subject to cost-effectiveness criteria (Castro et al., 2019; Soares, 2019). Courts, however, have adopted a strict literal interpretation of Article 196 ("health is a right of all and a duty of the state"), under which any individual can obtain virtually any medication through litigation—from common analgesics to rare-disease treatments costing over US$400,000 per year. Access is inexpensive (a prescription and a public defender suffice); approximately 85% of first-instance claims are granted (CNJ/INSPER, 2019). The result is 96,000 health court orders nationally in 2017, a 913% increase in São Paulo since 2008 (Online Appendix, Figure A.1). Litigated health spending in São Paulo alone amounts to approximately US$300 million annually—nearly 5% of the state's public health budget.

2.2. The Sanction Channel: Administrative vs. Litigated Purchases

The São Paulo State Department of Health (SES/SP) manages health policy through 99 decentralized public buyer units (PBUs). All purchases—ordinary and urgent—are conducted through the *Bolsa Eletrônica de Compras* (BEC), São Paulo's electronic procurement platform, mandatory for common goods since 2007. Two competitive procedures are used: sealed-bid tendering (*convite*) and multi-round descending auctions (*pregão*). For ordinary purchases, the internal phase typically spans 30–180 days; almost all court decisions (99.94%) are enforced as preliminary injunctions with delivery deadlines of 1–10 days, compressing the internal phase to roughly one-third of the ordinary timeline.

A critical institutional feature for our analysis is the distinction between *litigated* and *administrative* urgent purchases. Both types face identical planning constraints—compressed timelines, small quantities, diverted budget resources, and the same BEC platform and bidding procedures. They differ sharply in the consequences of procurement failure. For litigated purchases, failure to comply with a court order exposes public officials to substantial fines (sometimes disproportionate to the purchase value), administrative and criminal liability, and seizure or blocking of public funds; these penalties are public information, as the court order must be referenced in the tender notice. Since 2009, the SES/SP has run a parallel *administrative-request* channel: a mechanism allowing the government to negotiate supply directly with individuals before litigation occurs. Administrative requests undergo evaluation by a scientific committee using evidence-based medicine criteria, and—crucially—failure to complete an administrative purchase carries no penalties for public officials. Between 2009 and 2019, administrative requests generated approximately 9,700 purchases, or about 6.5% of the total from court orders.

This channel structure is the institutional vehicle for our identification. Administrative and litigated urgent purchases share all execution constraints but differ only in the penalty regime, so a within-item comparison between them is the cleanest available measurement of the cost margin attributable to personal sanctions on procurement officials. Figure 1 summarizes the three purchase types. We discuss the selection of cases into the administrative channel—and the threats it poses to a strict causal interpretation—in Section 4.

Table 1: Purchase Types: Institutional Characteristics

	Ordinary	Administrative	Litigated
Source of funds	Dedicated budget	Diverted budget	Diverted budget
Quantity	Large	Small	Small
Delivery time	Long (30–180 d)	Short (1–10 d)	Short (1–10 d)
Threat of punishment	None	None	**Fines, liability, fund seizure**

Notes: Administrative and litigated purchases share planning constraints (short timelines, small quantities, diverted budgets) and differ only in the penalty regime. The "under the gun" comparison exploits this variation.

3. Data and Sample

3.1. Data Sources

We combine three administrative datasets covering January 2009 through December 2019: (i) bid-level records from the BEC electronic procurement platform for all SES/SP purchases of medical supplies (BEC Group 65), where each observation is a purchase-offer-item (POI) recording reference prices, bid prices, number of bidding firms, quantities, and tender outcomes; (ii) purchase-type classification, assigning each POI as ordinary, administrative, or litigated using a regular-expression algorithm applied to tender notices, which by law must reference any underlying court order;[1] and (iii) health litigation records from S-CODES, the SES/SP database of health-related lawsuits, from which we derive SUS formulary status and injunction enforcement data. The binary outcome *tender success* equals one if the POI closes with an accepted winning bid within its procurement order and zero otherwise; re-tendered items enter the data as new POIs and are not counted retroactively as successes. All continuous variables are winsorized at the 1st and 99th percentiles; robustness to alternative winsorization is in the Appendix.

3.2. Sample Construction

The raw dataset comprises 479,330 POI observations across 33,144 distinct items, 51,059 purchase orders, and 100 PBUs (descriptive statistics in Online Appendix, Table A.1).

The analysis sample is restricted to items for which at least one ordinary and at least one litigated purchase is observed, ensuring within-item comparability. This yields 226,305 observations covering 3,856 items.[2] For regressions using bid prices, we further restrict to winning bids (196,995 ob-

[1] Online Appendix §A.7 documents the classifier's design and a stratified hand-labeling protocol over 500 notices (∼167 per predicted class). The regex relies on literal court-order references that are present or absent independently of the price outcomes we estimate, so any residual misclassification attenuates rather than biases our estimates.

[2] Sample sizes vary slightly across tables due to missing values in specific dependent variables and the restriction to winners for price regressions.

servations).

Several features of the data are noteworthy. In levels, ordinary purchases have substantially higher mean reference prices (R$909) compared to administrative (R$226) and litigated (R$370) purchases, reflecting the mix of items across procurement types. However, the composition effect is absorbed by our item fixed effects. In logs—which form the basis of our regressions—litigated purchases show higher mean prices (1.56 for reference price, 1.02 for negotiated price) than ordinary purchases (0.93 and 0.32), consistent with a price premium conditional on the item. Quantities are markedly lower for litigated purchases (mean of 9,646 units vs. 31,487 for ordinary), reflecting the inability to aggregate demand under compressed timelines. The number of bidding firms is also lower for urgent purchases (2.2 for litigated vs. 3.2 for ordinary), suggesting reduced competition in the external phase.

The Online Appendix presents complementary geographic evidence on the spatial distribution of administrative demand, which mirrors the patterns documented for litigation in Section 2.

For the "under the gun" analysis, we construct a subsample of urgent purchases (administrative and litigated) for items with both types present, yielding 56,803 observations (balance statistics in Online Appendix, Table A.2).

4. Empirical Strategy

The analysis has two stages: we first estimate the overall cost of urgency by comparing ordinary and urgent purchases, then isolate the sanction channel by comparing litigated and administrative urgent purchases.

4.1. Identification

Identification requires that, conditional on fixed effects, urgency is uncorrelated with unobserved determinants of procurement outcomes. Three institutional features support this. First, court orders originate from patients' health needs and legal claims—not from the procurement process. Brazilian public administration operates under a principle of impersonality: tender outcomes depend on item specifications and market conditions, not on who requested the purchase. Second, court orders are unpredictable from the buyer's perspective. Only 4% of claims arise from stock shortages or service failures, so litigation is poorly correlated with PBU-level mismanagement. Third, the urgency regime is determined externally—by judges or by the SES/SP's scientific committee—not by the purchasing unit. While PBUs could in principle misclassify purchases, the legal requirement to reference court orders in tender notices, combined with our REGEX-based classification of official texts, limits this concern. An honest-DiD event study around each item's first court order, computed via the imputation estimator of Borusyak et al. (2024) (Online Appendix, Figure A.15), exhibits pre-period coefficients statistically and economically indistinguishable from zero—supporting parallel trends—and post-period coefficients that rise monotonically from +5.4% at $t = 0$ to +15.8% at $t = +5$. This pattern rules out the reverse-causality story in which high prices attract lawsuits.

Selection into the administrative channel.. The under-the-gun comparison rests on the institutional fact that administrative and litigated urgent purchases share planning constraints and differ only in penalty exposure. Cases enter the administrative channel rather than the courts when the patient (or the patient's caregiver, often via SUS social workers) requests the medication directly through an SES/SP intake, and the request is approved by a scientific committee using cost-effectiveness criteria. Cases that bypass the administrative channel either (a) were not submitted to it because the patient went straight to litigation, or (b) were rejected by the committee. The committee thus selects on observed and unobserved item characteristics (essentiality, drug class, expected cost, evidence of clinical benefit). The direction of the resulting bias on our UTG coefficient is theoretically ambiguous: if the committee admits cost-effective items and rejects expensive outliers, our administrative subsample under-states the average cost of urgent procurement and our UTG gap is biased upward; if the committee admits items that are simpler to source and rejects clinically complex ones, our administrative subsample under-states the difficulty of the purchase and our UTG gap is biased downward. We do not attempt to sign this bias and instead frame the UTG estimate as a within-regime decomposition rather than a structural counterfactual against a hypothetical sanction-free regime; the latter object is unidentified in our setting.

Within-item balance.. Item fixed effects absorb time-invariant item characteristics but not within-item composition shifts: if severity, dose, or packaging systematically differ across purchase types, our estimates conflate procurement-channel effects with case mix. Within-item balance on pre-determined covariates (Online Appendix, Table A.3) supports the identifying assumption on observables: SUS-formulary status, time period, and buyer size are jointly balanced ($p > 0.26$ for each). The two dimensions that differ within item—log reference price and log quantity—are precisely the outcomes the paper attributes to the sanction mechanism (officials set looser reference prices and place smaller orders under judicial pressure); their imbalance is a *feature* of the treatment, not a confounder. Procurement modality (sealed-bid vs. reverse auction) differs by 4 pp, a gap small relative to the 23–30% UTG premium and absorbed by year-month fixed effects.

Other threats.. Three further threats warrant discussion. *First*, court orders originate from patients' health needs and legal claims, not from the procurement process. Brazilian public administration operates under a principle of impersonality: tender outcomes depend on item specifications and market conditions, not on who requested the purchase. Only 4% of claims arise from stock shortages or service failures, so litigation is poorly correlated with PBU-level mismanagement. *Second*, an honest-DiD event study around each item's first court order, computed via the imputation estimator of Borusyak et al. (2024) (Online Appendix, Figure A.15), exhibits pre-period coefficients economically and statistically indistinguishable from zero—supporting parallel trends—and post-period coefficients that rise from +5.4% at $t = 0$ to +15.8% at $t = +5$. The dynamic event-study evidence rules out the reverse-causality story in which high prices attract lawsuits and complements the cross-sectional fixed-effects design. *Third*, regex-based classification of tender notices may contain asymmetric misclassification between administrative and litigated requests; the classifier's design and hand-labeling protocol are documented in Online Appendix §A.7. Together with the never-litigated placebo (Section 5.5), the design narrows the residual identification space without eliminating it. Throughout the paper we frame the UTG estimate as a descriptive decomposition rather than a structural parameter; this is the central interpretive claim, and we honor it by reporting an explicit channel decomposition rather than a single counterfactual price.

4.2. Enforcement Costs: Ordinary vs. Urgent

We estimate the following baseline specification:

$$y_{i,g,t} = \beta \cdot \text{Urgent}_{i,g,t} + \gamma_g + \lambda_t + \delta_b + \varepsilon_{i,g,t}, \tag{1}$$

where $y_{i,g,t}$ is the outcome for purchase order i of item g at time t; γ_g are item fixed effects; λ_t are time fixed effects; and δ_b are PBU fixed effects. Standard errors are clustered at the PBU level. We report four specifications with progressively richer fixed effects, from item FE only through item + year-month + PBU FE. Our preferred specification (item + year + PBU FE) absorbs time-invariant item characteristics, common annual trends, and buyer-specific heterogeneity, while retaining sufficient within-cell variation for precise estimation.

For negotiated prices and number of firms, we additionally estimate a specification that partials out log quantity:

$$y_{i,g,t} = \beta \cdot \text{Urgent}_{i,g,t} + \alpha \cdot \ln Q_{i,g,t} + \gamma_g + \lambda_t + \delta_b + \varepsilon_{i,g,t}, \tag{2}$$

Because log quantity is itself an outcome of urgency, (2) does not identify a structural direct effect (Acharya et al., 2016); we therefore read the contrast between (1) and (2) as a descriptive mediation decomposition, quantifying how much of the total premium co-moves with the quantity margin that urgency itself induces. Section 5.1 develops this interpretation.

4.3. The "Under the Gun" Effect: Litigated vs. Administrative

To isolate the effect of judicial sanctions, we restrict the sample to urgent purchases only and estimate:

$$\ln(\text{Price}_{i,g,t}) = \beta \cdot \text{Admin}_{i,g,t} + \gamma_g + \lambda_t + \delta_b + \varepsilon_{i,g,t}, \tag{3}$$

where Admin equals 1 for administrative purchases and 0 for litigated. Since both types face identical planning constraints, β isolates the cost attributable to the sanction regime.

5. Results

5.1. Enforcement Costs: Prices, Competition, and Tender Success

Urgency distorts every stage of the procurement process. Reference prices—set during the planning phase—are 2.7% higher in the preferred specification ($p < 0.10$; Online Appendix, Table A.4). The raw premium is 17.8% with item FE alone but attenuates sharply with PBU FE, indicating that buyer-level heterogeneity accounts for much of the unconditional gap.

Negotiated prices are the core outcome (Table 2). Panel A reports the total effect; Panel B controls for log quantity.

Total effect (Panel A).. Negotiated prices are significantly higher for urgent purchases across all specifications. The total effect ranges from 17.3% ($e^{0.159} - 1$) with item FE only to 5.4–5.7% with our preferred specifications (columns 3–4). The attenuation pattern mirrors that of reference prices, confirming the importance of buyer-level heterogeneity.

Mediation decomposition (Panel B).. Adding log quantity as a control reduces the coefficient on urgency by roughly half in the most saturated specifications, from 0.053 to 0.030 (column 3). Log quantity is a post-treatment variable, so conditioning on it re-weights the comparison toward observations with similar order size (Acharya et al., 2016). The contrast between Panels A and B is therefore a within-paper mediation decomposition: of the 5.4% total premium, roughly half co-moves with quantity differences induced by urgency, while a residual 3.0% ($e^{0.030} - 1$, $p < 0.10$) remains once order-size

Table 2: Negotiated Prices

	(1)	(2)	(3)	(4)
Panel A: Total Effect				
Urgent Purchase	0.159***	0.151***	0.053***	0.056***
	(0.042)	(0.038)	(0.016)	(0.016)
Panel B: Direct Effect				
Urgent Purchase	0.075	0.060	0.030*	0.030*
	(0.045)	(0.042)	(0.017)	(0.017)
Log Quantity	-0.321***	-0.325***	-0.392***	-0.393***
	(0.026)	(0.026)	(0.029)	(0.029)
Item FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	No
Year-Month FE	No	No	No	Yes
PBU FE	No	No	Yes	Yes
Observations	196,886	196,886	196,883	196,883
Within R^2 (A)	0.003	0.003	0.000	0.000
Within R^2 (B)	0.315	0.327	0.358	0.359

Notes: Dependent variable: log negotiated price. Sample: winners only. Panel A reports the total effect; Panel B controls for log quantity. Winsorized at 1%/99%. Standard errors clustered at the PBU level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

variation is partialled out. The coefficient on log quantity is large and negative (-0.392), consistent with the bulk-discount channel that is the central piece of the demand-fragmentation mechanism developed below.

5.2. Competition and Tender Success

Urgent purchases attract -5.4% fewer bidding firms (-0.056, $p < 0.01$; Online Appendix, Table A.6), consistent with reduced competition under compressed timelines (Bulow and Klemperer, 1996). Quantities also fall, though imprecisely (-0.058, not significant; Online Appendix, Table A.5).

Yet urgent tenders are 2.1 pp *more* likely to succeed ($p < 0.01$; Online Appendix, Table A.7)—precisely what the "under the gun" mechanism predicts: officials accept worse terms to ensure the purchase goes through. This also rules out sample selection: urgent tenders succeed *despite* less favorable conditions.

5.3. The "Under the Gun" Effect

The sharpest comparison contrasts litigated and administrative urgent purchases (Table 3). Both share the compressed internal phase, the small-order, diverted-budget regime of urgent procurement; they differ only in whether the officer faces personal fines and fund seizure if the tender fails. We frame the resulting gap as a within-regime decomposition: it tells us how much cost margin is associated with the sanction-exposed regime relative to the sanction-free regime, holding planning constraints fixed.

Total effect (Panel A).. Administrative purchases have significantly lower negotiated prices than litigated ones across all specifications. The implied litigated premium ranges from 23.2% with item fixed effects alone (column 1) to 30.4% with item and year-month fixed effects plus PBU controls (column 4). In our preferred specification (column 3), the coefficient is -0.262 ($p < 0.05$), implying that litigated purchases are 30.0% more expensive than comparable administrative purchases. The same-item, same-month comparison (column 4) yields -0.265 on $N = 27{,}122$ within-cell observations, ruling

Table 3: Under the Gun: Administrative vs Litigated

	(1)	(2)	(3)	(4)
Panel A: Total Effect				
Administrative (vs Litigated)	-0.209***	-0.205**	-0.262**	-0.265***
	(0.074)	(0.074)	(0.096)	(0.094)
Panel B: Direct Effect				
Administrative (vs Litigated)	-0.006	0.016	0.117	0.115
	(0.056)	(0.058)	(0.122)	(0.115)
Log Quantity	-0.284***	-0.292***	-0.341***	-0.344***
	(0.036)	(0.038)	(0.042)	(0.040)
Item FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	No
Year-Month FE	No	No	No	Yes
PBU FE	No	No	Yes	Yes
Observations	56,803	56,803	56,803	56,802
Within R^2 (A)	0.011	0.010	0.007	0.007
Within R^2 (B)	0.287	0.292	0.307	0.310

Notes: Dependent variable: log negotiated price. Sample: urgent purchases only (administrative and litigated), winners, items with both types present. Panel A reports the total effect; Panel B controls for log quantity. Winsorized at 1%/99%. Standard errors clustered at the PBU level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

out time-varying within-item confounders as a driver.[3]

Quantity channel does the work (Panel B).. Adding log quantity as a control pulls the UTG coefficient close to zero (0.117, SE = 0.122, column 3); under the tightest item×year-month FE specification, the coefficient is statistically indistinguishable from zero (0.136, SE = 0.122). Within UTG,

[3]Item×year-month FE absorb 30,029 singleton observations, leaving 27,122 with within-cell variation. Of 40,966 item×month cells, 4,402 (11%) contain both purchase types.

the litigated-vs-administrative gap is essentially mediated by order quantity: holding quantity constant, sanctioned and unsanctioned urgent purchases of the same item in the same month face the same per-unit price.

The magnitude of the quantity gap is large enough to mechanically generate a price gap as big as the one we observe. The log-quantity gap between administrative and litigated purchases (1.34) times the bulk-discount elasticity (-0.34) accounts for a price gap of about 58 from fragmentation alone—of the same order of magnitude as the 30.0% total, and slightly larger. Equivalently, regressing log quantity on the administrative indicator under our preferred specification yields a coefficient of 1.203: for the same item, administrative orders are roughly 3.3 times the size of litigated orders. Sanctions do not merely correlate with smaller orders; they shrink demand aggregation by a multiple sufficient to mechanically explain the entire UTG price gap through the bulk-discount channel. We return to this point in Section 5.4, where we decompose the urgency premium and the UTG gap into their constituent channels.

5.4. Three-Channel Decomposition

The cost margin between regimes is an umbrella that we open into three distinct channels: *demand fragmentation* (C1), *reduced participation* (C2), and *supplier composition shift* (C3). To isolate them, we run the same urgency regression under four progressively richer specifications: (1) a baseline with item, year, and PBU fixed effects; (2) the baseline plus log quantity;

(3) the baseline plus firm fixed effects; and (4) the baseline plus both log quantity and firm fixed effects. Comparing across specifications attributes the drop from (1) to (2) to the quantity channel, the drop from (1) to (3) to the firm-selection channel, and the residual in (4) to the participation channel and any within-firm markup.

Table 4 reports the cascade for the urgent-vs-ordinary comparison; Table 5 for the UTG comparison.

Table 4: Three-Channel Decomposition: Urgent vs. Ordinary Procurement

	(1) Baseline	(2) + log qty.	(3) + firm FE	(4) + qty + firm FE
Urgent purchase	0.051***	0.028	0.024	0.011
	(0.015)	(0.018)	(0.016)	(0.018)
Implied premium (%)	5.26%	2.85%	2.48%	1.10%
Item FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
PBU FE	Yes	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
log Quantity control	No	Yes	No	Yes
Observations	196,883	196,883	196,642	196,642

Notes: Three-channel decomposition. Column (1) gives the total premium under the preferred specification. Columns (2)–(4) progressively absorb channels: (2) absorbs *C1 demand fragmentation* via log quantity; (3) absorbs *C3 supplier composition shift* via firm fixed effects; (4) absorbs both, leaving the C2 *competition channel* residual (and any within-firm markup). Decomposition: total 5.26% = 2.41 pp (C1 quantity) + 2.79 pp (C3 firm selection) + 1.10 pp (residual). Standard errors clustered at the PBU level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Across the full sample.. The 5.26% urgent-vs-ordinary cost margin decomposes into 2.41 pp from the quantity channel (C1), 2.79 pp from the firm-selection channel (C3), and a 1.10 pp residual shared by the C2 participa-

Table 5: Three-Channel Decomposition: Litigated vs. Administrative (Under the Gun)

	(1) Baseline	(2) + log qty.	(3) + firm FE	(4) + qty + firm FE
Administrative (vs Lit.)	-0.258***	0.138	-0.184***	0.157
	(0.094)	(0.129)	(0.066)	(0.119)
Implied premium (%)	29.44%	-12.88%	20.14%	-14.54%
Item FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
PBU FE	Yes	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
log Quantity control	No	Yes	No	Yes
Observations	57,084	57,084	56,977	56,977

Notes: Three-channel decomposition. Column (1) gives the total premium under the preferred specification. Columns (2)–(4) progressively absorb channels: (2) absorbs *C1 demand fragmentation* via log quantity; (3) absorbs *C3 supplier composition shift* via firm fixed effects; (4) absorbs both, leaving the C2 *competition channel* residual (and any within-firm markup). Decomposition: total 29.44% = 42.32 pp (C1 quantity) + 9.29 pp (C3 firm selection) + -14.54 pp (residual). Standard errors clustered at the PBU level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

tion channel and any within-firm markup. C1 captures the bulk-discount loss from sanctions-induced fragmentation. C3 captures the equilibrium consequence of urgency on which firm wins: the supplier mix shifts toward compressed-timeline specialists, systematically different from—though not necessarily more expensive than—suppliers serving ordinary procurement. C2 captures the participation-margin compression: at identical order size, urgent tenders attract -4.2% fewer bidders (-0.043, SE 0.012, $p < 0.001$; §5.2), thinning the bid distribution and raising negotiated prices.

Within the under-the-gun comparison.. The UTG decomposition (Table 5) is sharper: of the 29.44% total, the quantity channel absorbs 42.32 pp—more than the entire premium—and the residual after both controls is -14.54 pp, indistinguishable from zero. Within UTG, the cost margin is *entirely* mediated by demand fragmentation; personal sanctions on procurement officials, conditional on urgency, generate no per-unit-price differential beyond what the quantity channel mechanically explains.

Reference prices reinforce the channel reading.. A separate diagnostic confirms that the C2 participation channel, not officials' price-setting behavior, is what generates the residual urgency premium. Re-estimating the Panel B specification with reference prices (set *before* bidding, so unaffected by supplier markup or selection) yields a residual coefficient of 0.004 (SE 0.024), statistically zero. Officials under urgency do not set looser reference prices; the residual urgency premium on negotiated prices arises after bidding, in thinner markets.

Item event study confirms dynamic fragmentation.. Within the same item, the arrival of the first court order causes order quantity to drop sharply: -27.8% at the year of treatment, recovering to -0.3% by year +5 (Online Appendix, Figure A.16). The price event-study reported earlier moves in the opposite direction; the joint pattern is consistent with the fragmentation mechanism operating in real time at the item level.

5.5. Falsification: Items Never Litigated

If the urgency premium reflects the causal effect of judicial pressure rather than unobserved item characteristics, it should be absent for items that are never subject to court orders. Running the preferred specification on items with zero litigated purchases across 2009–2019 (restricted to those with both ordinary and administrative purchases) yields placebo coefficients of -0.020 (SE $= 0.032$) for negotiated prices and -0.031 (SE $= 0.045$) for reference prices—both economically small and statistically insignificant (Online Appendix, Table A.8). The urgency premium exists only for items that are targets of litigation, ruling out a generic correlation between urgency and prices.

5.6. Supplier Composition Shift, Not Within-Firm Markup

The C3 channel in Table 4 reflects equilibrium changes in *which* firms win urgent contracts, not the pricing behavior of any one firm. We confirm this directly by estimating the urgency coefficient within firm-buyer-item triples observed under both ordinary and urgent procurement—the tightest possible specification. Of the 8,232 such triples (covering 33,186 observations), the urgency coefficient is -0.027 (SE 0.010, $p < 0.01$): same firm, same buyer, same item, the price under urgency is *lower* by -2.67%, not higher. The within-firm-markup interpretation of supplier fixed effects—an intuition often invoked in the supply-side under-the-gun literature—is not what we find.

Two facts support the selection reading. First, 92% of urgent winners also serve ordinary procurement: the supplier pool is largely shared, so selection operates at the equilibrium-weighting margin rather than through entry/exit.

Second, the channel is most pronounced in less concentrated markets (8.5% firm-FE residual) and absent in concentrated ones (-2.0%), the opposite of what a within-firm markup would predict. The fiscal cost C3 carries is a sourcing problem, not a price-discrimination problem: under urgency, fewer specialist firms are available and the equilibrium falls on a different point of the supplier distribution.

5.7. Where the Cost Bites Hardest

Splitting items by mean number of participating firms (median 2.071), competitive items show an 14.4% urgency premium with 8.5% surviving firm fixed effects (39 attenuation). Concentrated items show a much smaller, statistically insignificant baseline. The pattern matches the channel logic: urgency erodes competitive pressure where there is competitive pressure to erode. Where there is none—monopolistic items where the buyer has no outside option in any case—urgency carries little additional cost.

5.8. Summary of Effects

Figure 1 summarizes the three-channel decomposition for both comparison levels. Across the urgent-vs-ordinary comparison, the 5.26% premium splits into 2.41 pp from C1 (quantity), 2.79 pp from C3 (supplier selection), and a 1.10 pp residual (C2 participation plus any within-firm markup). Within the under-the-gun comparison, the 29.44% gap is fully mediated by the quantity channel; the residual is statistically zero.

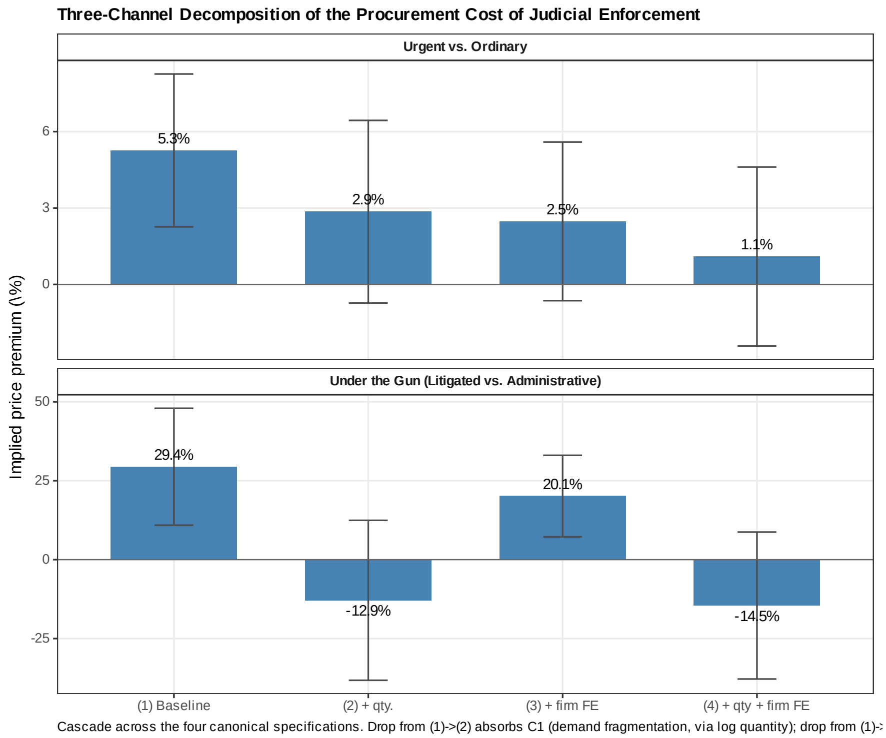

Figure 1: Three-Channel Decomposition of the Procurement Cost of Judicial Enforcement
Notes: Cascade across the four canonical specifications. Drop from (1) to (2) absorbs C1 (demand fragmentation, via log quantity); drop from (1) to (3) absorbs C3 (supplier composition shift, via firm fixed effects); column (4) is the residual. Top panel: urgent vs. ordinary across the full sample. Bottom panel: litigated vs. administrative within the urgent sub-sample (UTG). All specifications include item, year, and PBU fixed effects; standard errors clustered at the PBU level. Error bars: 95% confidence intervals.

6. Conclusion

Courts that compel governments to buy medications save individual lives. They also cost the public budget 5.4% more per purchase across the full sample, attract -5.4% fewer bidders, and—when the procurement officer is exposed to personal sanctions—generate a 23–30% cost margin between litigated and administrative urgent purchases of the same item.

The decomposition we develop traces this margin to three distinct channels operating simultaneously. *Demand fragmentation* (C1) does most of the work: sanctions force smaller orders, and the bulk-discount channel mechanically explains essentially all of the under-the-gun gap and roughly half of the urgency premium across the full sample. *Reduced participation* (C2) accounts for most of the across-sample residual: at identical order size, urgent tenders attract -4.2% fewer bidders, leaving thinner markets and worse negotiated prices. *Supplier composition shift* (C3) accounts for the rest of the across-sample premium: under urgency, the supplier mix shifts toward firms that specialize in compressed-timeline delivery, yielding worse matches on average without any within-firm markup. The within firm-buyer-item triple test directly rules out the within-firm-markup interpretation that supplier fixed effects might suggest in isolation: the same firm sells the same item to the same buyer at a slightly *lower* per-unit price under urgency, not a higher one.

The decomposition has direct policy content. Two complementary reforms address two distinct channels.

On the demand side, expanding São Paulo's administrative-request mechanism—which since 2009 has allowed urgent procurement without personal-sanction exposure—would eliminate the channel that operates inside UTG. Applied to a defensible target of 50% of court-mandated spending, this would recover on the order of $44 M per year on the 300 M of São Paulo's annual litigated spending. The instrument is administrative capacity (the scientific-committee infrastructure), not legal reform.

On the supply side, framework agreements for items where fragmentation is the binding cost—repeat-purchase commodity-like medications—directly address C1 by allowing aggregation across urgent and ordinary requests. Calibrated to a feasible 40% of the annual spending, the C1 channel mapped onto this scope would recover on the order of \$3 M per year. Limiting the prominence of court-order references in tender notices (which currently serves as a public signal of the government's reduced outside option) would address C3 without constraining the underlying litigation; we view this as second-order relative to the C1 reform.

The fiscal scale of the cost is substantial but bounded. Applying the urgent-vs-ordinary premium to the \$300 M of annual litigated spending in São Paulo gives a per-unit-price-margin lower bound of \$16 M per year. Applying the under-the-gun premium to the \$250–300 M of sanction-exposed spending gives an upper bound of \$74–88 M per year. The publishable range is \$16–88 M per year on the per-unit-price margin alone; adding the welfare cost of smaller order sizes, reduced competition, and officials' search time diverted from ordinary procurement would raise this further.[4]

These findings highlight a tension at the heart of right-to-health enforcement: judicial mandates intended to guarantee individual access impose sub-

[4]The lower bound applies the headline urgent-vs-ordinary premium (5.4%) to the entire \$300 M base. The upper bound applies the under-the-gun premium (23–30%) to the sanction-exposed portion (\$250–300 M). Both are per-unit-price-margin estimates and exclude welfare losses from reduced competition and constrained order sizes.

stantial costs on the public budget, potentially reducing resources available for the broader health system. Institutional designs that balance individual rights with procurement efficiency—rather than treating them as competing objectives—could improve welfare for both litigants and the broader population. Brazil's recent procurement reform (Lei 14.133/2021), which expands electronic auctions and introduces framework agreements, could mitigate the quantity-fragmentation channel we document—a prediction future work can test as the reform takes effect.

References

Acharya, A., Blackwell, M., Sen, M., 2016. Explaining causal findings without bias: Detecting and assessing direct effects. American Political Science Review 110, 512–529. doi:10.1017/S0003055416000216.

Baltrunaite, A., Giorgiantonio, C., Mocetti, S., Orlando, T., 2021. Discretion and supplier selection in public procurement. Journal of Law, Economics, and Organization 37, 134–166. doi:10.1093/jleo/ewaa009.

Bandiera, O., Best, M.C., Khan, A.Q., Prat, A., 2021. The allocation of authority in organizations: A field experiment with bureaucrats. Quarterly Journal of Economics 136, 2195–2242. doi:10.1093/qje/qjab029.

Bandiera, O., Prat, A., Valletti, T., 2009. Active and passive waste in government spending: Evidence from a policy experiment. American Economic Review 99, 1278–1308. doi:10.1257/aer.99.4.1278.

Best, M.C., Hjort, J., Szakonyi, D., 2023. Individuals and organizations as sources of state effectiveness. American Economic Review 113, 2121–2167. doi:10.1257/aer.20191598.

Biehl, J.a., Petryna, A., Gertner, A., Amon, J.J., Picon, P.D., 2009. Judicialisation of the right to health in Brazil. The Lancet 373, 2182–2184. doi:10.1016/S0140-6736(09)61172-7.

Borusyak, K., Jaravel, X., Spiess, J., 2024. Revisiting event-study designs: Robust and efficient estimation. Review of Economic Studies 91, 3253–3285. doi:10.1093/restud/rdae007.

Bulow, J., Klemperer, P., 1996. Auctions versus negotiations. American Economic Review 86, 180–194.

Callaway, B., Sant'Anna, P.H.C., 2021. Difference-in-differences with multiple time periods. Journal of Econometrics 225, 200–230. doi:10.1016/j.jeconom.2020.12.001.

Carril, R., Gonzalez-Lira, A., Walker, M.S., 2026. Competition under incomplete contracts and the design of procurement policies. American Economic Review 116, 535–581. doi:10.1257/aer.20221345.

Castro, M.C., Massuda, A., Almeida, G., Menezes-Filho, N.A., Andrade, M.V., de Souza Noronha, K.V.M., Rocha, R., Macinko, J., Hone, T., Tasca, R., Giovanella, L., Malik, A.M., Werneck, H., Fachini, L.A., Atun, R., 2019. Brazil's unified health system: The first 30 years and prospects for the future. The Lancet 394, 345–356. doi:10.1016/S0140-6736(19)31243-7.

CNJ/INSPER, 2019. Judicialização da saúde no Brasil: Perfil das demandas, causas e propostas de solução. Conselho Nacional de Justiça, Brasília.

Coviello, D., Guglielmo, A., Spagnolo, G., 2018a. The effect of discretion on procurement performance. Management Science 64, 715–738. doi:10.1287/mnsc.2016.2628.

Coviello, D., Moretti, L., Spagnolo, G., Valbonesi, P., 2018b. Court efficiency and procurement performance. Scandinavian Journal of Economics 120, 826–858. doi:10.1111/sjoe.12225.

Decarolis, F., Giuffrida, L.M., Iossa, E., Mollisi, V., Spagnolo, G., 2020. Bureaucratic competence and procurement outcomes. Journal of Law, Economics, and Organization 36, 537–597. doi:10.1093/jleo/ewaa004.

Ferraz, O.L.M., 2009. The right to health in the courts of Brazil: Worsening health inequities? Health and Human Rights 11, 33–45.

Glaeser, E.L., Shleifer, A., 2003. The rise of the regulatory state. Journal of Economic Literature 41, 401–425. doi:10.1257/002205103765762725.

Lewis-Faupel, S., Neggers, Y., Olken, B.A., Pande, R., 2016. Can electronic procurement improve infrastructure provision? Evidence from public works in India and Indonesia. American Economic Journal: Economic Policy 8, 258–283. doi:10.1257/pol.20140258.

Prendergast, C., 2007. The motivation and bias of bureaucrats. American Economic Review 97, 180–196. doi:10.1257/aer.97.1.180.

Rasul, I., Rogger, D., 2018. Management of bureaucrats and public service delivery: Evidence from the Nigerian civil service. Economic Journal 128, 413–446. doi:10.1111/ecoj.12418.

Soares, A., 2019. Health system financing paradigm in the state of São Paulo: A regional analysis. Revista de Saúde Pública 53, 39. doi:10.11606/S1518-8787.2019053000796.

Szücs, F., 2024. Discretion and favoritism in public procurement. Journal of the European Economic Association 22, 117–160. doi:10.1093/jeea/jvad030.

Wang, D.W.L., 2015. Right to health litigation in Brazil: The problem and the institutional responses. Human Rights Law Review 15, 617–641. doi:10.1093/hrlr/ngv025.

Williams, M.J., 2021. Beyond state capacity: Bureaucratic performance, policy implementation, and reform. Journal of Institutional Economics 17, 339–357. doi:10.1017/S1744137420000478.

Online Appendix

This appendix provides tables and figures from the main text, robustness checks, and supplementary evidence.

A.0 Supporting Descriptives, Balance, and Secondary Outcome Tables

Table A.1: Descriptive Statistics by Purchase Type

	Ordinary		Administrative		Litigated		Diff	Diff
	Mean	(SD)	Mean	(SD)	Mean	(SD)	(O–L)	(A–L)
Panel A: Levels								
Reference Price	909.17	(5,959.41)	226.38	(2,257.35)	370.39	(2,540.20)	538.78***	-144.01***
							[0.000]	[0.000]
Negotiated Price	477.51	(2,915.88)	148.03	(1,177.52)	260.25	(1,561.81)	217.26***	-112.22***
							[0.000]	[0.000]
Quantity	31,487.15	(163,507.29)	117,146.00	(356,962.37)	9,646.19	(82,683.10)	21,840.96***	107,499.81***
							[0.000]	[0.000]
N. Bidding Firms	3.17	(1.93)	2.43	(1.35)	2.20	(1.19)	0.97***	0.23***
							[0.000]	[0.000]
Panel B: Log Transformations								
Log Reference Price	0.926	(2.723)	0.909	(2.541)	1.557	(2.446)	-0.632***	-0.648***
							[0.000]	[0.000]
Log Negotiated Price	0.319	(2.862)	0.480	(2.604)	1.024	(2.568)	-0.705***	-0.544***
							[0.000]	[0.000]
Log Quantity	6.947	(2.678)	7.472	(3.160)	6.087	(2.083)	0.860***	1.385***
							[0.000]	[0.000]
Log N. Firms	0.984	(0.590)	0.747	(0.529)	0.665	(0.493)	0.319***	0.081***
							[0.000]	[0.000]
Panel C: Tender Characteristics								
Successful Tender (%)	0.859	(0.348)	0.869	(0.338)	0.908	(0.288)	-0.050***	-0.040***
							[0.000]	[0.000]
Observations	136,250		15,371		45,351			

Notes: Sample restricted to items with at least one ordinary and one litigated purchase. Winners only for price and quantity variables. Variables winsorized at 1%/99%. Stars on differences from Welch t-tests; p-values in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Section A.1 examines the sensitivity of the main estimates to the choice of winsorization level. Section A.2 assesses the robustness of the "under the gun" estimates to progressive inclusion of control variables. Section A.3 presents distributional evidence on key outcome variables and additional descriptive figures. Section A.4 presents geographic descriptive evidence on

Table A.2: Balance Table: Administrative vs Litigated (Urgent Purchases)

	Administrative		Litigated		Diff	p-value
	Mean	(SD)	Mean	(SD)	(A–L)	
Panel A: Procurement Outcomes						
Reference Price	226.38	(2,257.35)	232.56	(1,432.18)	-6.18	[0.752]
Negotiated Price	148.03	(1,177.52)	174.79	(1,052.56)	-26.75**	[0.013]
Quantity	117,146.00	(356,962.37)	9,306.87	(80,252.25)	107,839.13***	[0.000]
Log Reference Price	0.91	(2.54)	1.43	(2.32)	-0.52***	[0.000]
Log Negotiated Price	0.48	(2.60)	0.89	(2.44)	-0.41***	[0.000]
Log Quantity	7.47	(3.16)	6.18	(1.99)	1.29***	[0.000]
Panel B: Market Structure						
N. Bidding Firms	2.427	(1.351)	2.155	(1.120)	0.272***	[0.000]
Log N. Firms	0.747	(0.529)	0.650	(0.483)	0.097***	[0.000]
Successful Tender (%)	0.869	(0.338)	0.914	(0.280)	-0.045***	[0.000]
Panel C: Purchase Characteristics						
Electronic Auction	0.985	(0.123)	0.941	(0.236)	0.044***	[0.000]
Observations	15,371		41,491			

Notes: Sample restricted to urgent purchases (administrative and litigated) for items with both types present. Winners only for price/quantity. Variables winsorized at 1%/99%. p-values from Welch t-tests in brackets. *** p<0.01, ** p<0.05, * p<0.1.

administrative demand, complementing the litigation maps shown in Section 2. Section A.5 presents an event study around the first court order for each item. Section A.6 examines heterogeneity in the urgency premium along several dimensions.

Table A.3: Within-Item Balance: Administrative vs. Litigated

Covariate	Mean (Admin)	Mean (Lit)	Raw diff. (A–L)	Within-item diff. Coef	Within-item diff. (SE)
SUS basic (MEDICAMENTO flag)	0.718	0.892	-0.174	0.000^{NA}	(NaN)
Electronic auction (pregão)	0.984	0.941	0.043	0.041^{**}	(0.017)
Log quantity (bulk size signal)	7.642	6.180	1.462	0.866^{*}	(0.502)
Log reference price	0.835	1.401	-0.566	-0.339^{***}	(0.093)
Successful tender	0.869	0.914	-0.045	-0.024	(0.022)
Late period (year $\geq$ 2014)	0.565	0.623	-0.058	-0.045	(0.041)
Large PBU (above-median)	0.895	0.810	0.085	0.083	(0.073)
Observations			63,426		
Items			2,370		

Notes: Within-item balance between administrative ($Admin = 1$) and litigated ($Admin = 0$) urgent purchases, restricted to items with both types present. Raw difference is the unconditional mean gap. Within-item coefficient is β from $x_{i,g,t} = \beta\,\text{Admin}_{i,g,t} + \gamma_g + \varepsilon_{i,g,t}$, where γ_g is an item fixed effect and standard errors are clustered at the PBU level. A coefficient close to zero means that, within a given item, the administrative and litigated draws are balanced on the covariate. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.4: Reference Prices

	(1)	(2)	(3)	(4)
Urgent Purchase	0.164***	0.156***	0.027*	0.027*
	(0.056)	(0.050)	(0.014)	(0.014)
Item FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	No
Year-Month FE	No	No	No	Yes
PBU FE	No	No	Yes	Yes
Observations	196,899	196,899	196,896	196,896
Within R^2	0.003	0.003	0.000	0.000

Notes: Dependent variable: log reference price. Sample: winners only, items with both ordinary and litigated purchases. Winsorized at 1%/99%. Standard errors clustered at the PBU level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A.5: Quantities

	(1)	(2)	(3)	(4)
Urgent Purchase	-0.264	-0.279	-0.058	-0.065
	(0.183)	(0.181)	(0.056)	(0.058)
Item FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	No
Year-Month FE	No	No	No	Yes
PBU FE	No	No	Yes	Yes
Observations	196,899	196,899	196,896	196,896
Within R^2	0.003	0.003	0.000	0.000

Notes: Dependent variable: log quantity. Sample: winners only. Winsorized at 1%/99%. Standard errors clustered at the PBU level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

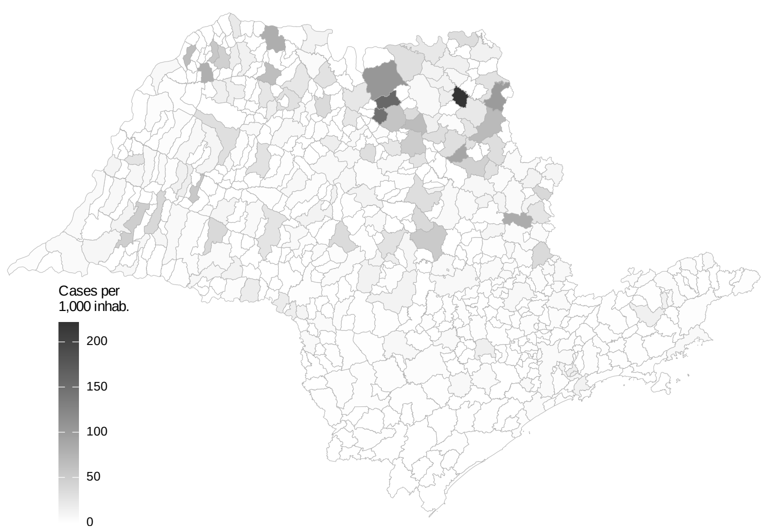

Figure A.1: Health Litigation Cases per 1,000 Inhabitants Across Municipalities in São Paulo (2009–2019)

Source: Authors' elaboration based on CNJ/INSPER (2019) litigation data and IBGE population estimates.

Table A.6: Participant Firms

	(1)	(2)	(3)	(4)
Panel A: Total Effect				
Urgent Purchase	-0.101***	-0.113***	-0.056***	-0.054***
	(0.023)	(0.022)	(0.014)	(0.015)
Panel B: Direct Effect				
Urgent Purchase	-0.091***	-0.102***	-0.042***	-0.040***
	(0.014)	(0.011)	(0.012)	(0.013)
Log Quantity	0.063***	0.064***	0.065***	0.065***
	(0.004)	(0.004)	(0.003)	(0.003)
Item FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	No
Year-Month FE	No	No	No	Yes
PBU FE	No	No	Yes	Yes
Obs. (Panel A)	81,805	81,805	81,801	81,801
Obs. (Panel B)	81,793	81,793	81,789	81,789
Within R^2 (A)	0.004	0.006	0.001	0.001
Within R^2 (B)	0.063	0.069	0.043	0.044

Notes: Dependent variable: log number of bidding firms. Sample: winners only. Panel A reports the total effect; Panel B controls for log quantity. Winsorized at 1%/99%. Standard errors clustered at the PBU level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A.7: Tender Success (LPM)

	(1)	(2)	(3)	(4)
Panel A: Total Effect				
Urgent Purchase	0.023**	0.023**	0.021***	0.021***
	(0.009)	(0.009)	(0.006)	(0.006)
Panel B: Direct Effect				
Urgent Purchase	0.023**	0.023**	0.021***	0.022***
	(0.009)	(0.010)	(0.006)	(0.006)
Log Quantity	0.003	0.003	0.012***	0.012***
	(0.004)	(0.004)	(0.002)	(0.002)
Item FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	No
Year-Month FE	No	No	No	Yes
PBU FE	No	No	Yes	Yes
Obs. (Panel A)	226,305	226,305	226,301	226,301
Obs. (Panel B)	226,282	226,282	226,278	226,278
Within R^2 (A)	0.001	0.001	0.000	0.000
Within R^2 (B)	0.001	0.001	0.004	0.004

Notes: Dependent variable: successful tender (LPM). Sample: all observations. Panel A reports the total effect; Panel B controls for log quantity. Winsorized at 1%/99%. Standard errors clustered at the PBU level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A.8: Placebo Test: Items Never Subject to Litigation

	bid_price_log		bid_price_ref_log	
	(1)	(2)	(3)	(4)
Urgent Purchase	-0.0204	0.0513***	-0.0306	0.0305**
	(0.0316)	(0.0153)	(0.0447)	(0.0141)
Observations	39,283	196,883	46,874	226,278
R^2	0.83364	0.86796	0.81121	0.85183
Within R^2	7.03×10^{-6}	0.00022	1.73×10^{-5}	7.43×10^{-5}
item_id fixed effects	✓	✓	✓	✓
year_n fixed effects	✓	✓	✓	✓
pbu_id fixed effects	✓	✓	✓	✓

Placebo sample: items with zero litigated purchases across 2009-2019. Main sample: items with at least one litigated and one ordinary purchase. DV: log price. Item + Year + PBU FE. SE clustered at PBU level.

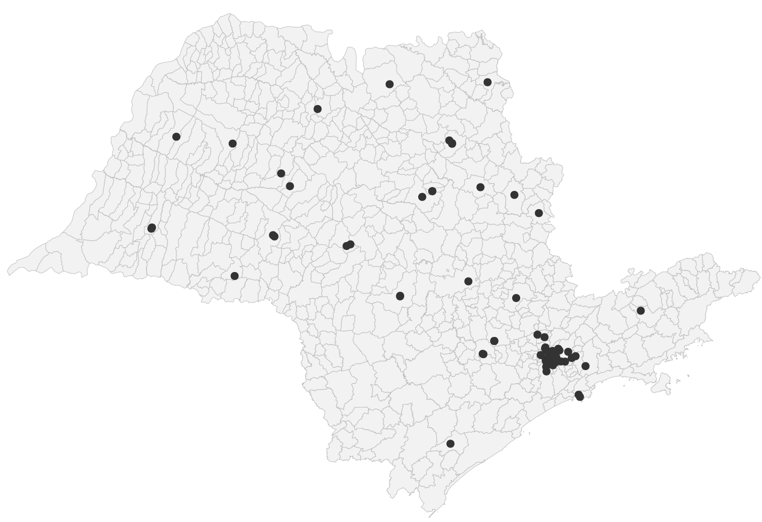

Figure A.2: Public Buyer Units (SES/SP)

Source: Authors' elaboration based on SES/SP administrative records.

A.1 Winsorization Sensitivity

Our baseline estimates use 1%/99% winsorization to limit the influence of extreme outliers while preserving meaningful variation in the data. To verify that our results are not driven by this specific choice, Tables A.9–A.12 replicate the main regression tables under three alternative winsorization treatments: no winsorization (Panel A), the 1%/99% baseline (Panel B), and a more aggressive 5%/95% winsorization (Panel C).

Table A.9 reports the sensitivity of the reference price estimates. The positive coefficient on urgent purchases is stable across all three winsorization levels: the preferred specification (column 3) yields estimates of 0.026 (no winsorization), 0.027 (baseline), and 0.027 with aggressive winsorization. The qualitative pattern—a large raw premium that attenuates substantially with PBU fixed effects—is preserved throughout, indicating that the reference price result is not driven by outliers.

Table A.10 presents the corresponding sensitivity analysis for negotiated prices. The total effect of urgency on negotiated prices is positive and statistically significant across all winsorization levels. The coefficient magnitudes are slightly larger without winsorization (reflecting the influence of extreme values) and slightly smaller with aggressive winsorization, but the economic message remains unchanged: urgent purchases are associated with higher negotiated prices, with the preferred estimate ranging from approximately 5% to 6%.

Table A.11 examines the robustness of the bidder participation results.

Table A.9: Robustness: Reference Prices

	(1)	(2)	(3)	(4)
Panel A: No Winsorization				
Urgent Purchase	0.167***	0.159***	0.026*	0.026*
	(0.056)	(0.051)	(0.014)	(0.014)
Panel B: Winsorized 1%/99%				
Urgent Purchase	0.164***	0.156***	0.027*	0.027*
	(0.056)	(0.050)	(0.014)	(0.014)
Panel C: Winsorized 5%/95%				
Urgent Purchase	0.156***	0.150***	0.027*	0.028**
	(0.053)	(0.048)	(0.014)	(0.014)
Item FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	No
Year-Month FE	No	No	No	Yes
PBU FE	No	No	Yes	Yes
Observations	196,899	196,899	196,896	196,896

Notes: Dependent variable: log reference price. Sample: winners only (success uses all obs). Each panel uses a different winsorization level. Standard errors clustered at the PBU level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

The negative effect of urgency on the number of bidding firms is remarkably stable across winsorization levels, with preferred estimates in the range of -0.06 to -0.05 log points. This stability is expected, since the number of firms is a count variable with limited extreme values, making it less sensitive to winsorization.

Table A.12 replicates the tender success analysis. The positive coefficient on urgency—indicating that urgent tenders are more likely to succeed—is robust across all winsorization levels. The magnitude is consistent at approximately 2 pp in the preferred specification, reinforcing the interpretation

Table A.10: Robustness: Negotiated Prices

	(1)	(2)	(3)	(4)
Panel A: No Winsorization				
Urgent Purchase	0.160***	0.152***	0.051***	0.054***
	(0.043)	(0.038)	(0.015)	(0.016)
Panel B: Winsorized 1%/99%				
Urgent Purchase	0.159***	0.151***	0.053***	0.056***
	(0.042)	(0.038)	(0.016)	(0.016)
Panel C: Winsorized 5%/95%				
Urgent Purchase	0.162***	0.153***	0.058***	0.061***
	(0.041)	(0.037)	(0.017)	(0.017)
Item FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	No
Year-Month FE	No	No	No	Yes
PBU FE	No	No	Yes	Yes
Observations	196,886	196,886	196,883	196,883

Notes: Dependent variable: log negotiated price. Sample: winners only (success uses all obs). Each panel uses a different winsorization level. Standard errors clustered at the PBU level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

that procurement officials under pressure accept less favorable terms to ensure compliance.

Table A.11: Robustness: Participant Firms

	(1)	(2)	(3)	(4)
Panel A: No Winsorization				
Urgent Purchase	-0.101***	-0.113***	-0.056***	-0.054***
	(0.023)	(0.022)	(0.015)	(0.015)
Panel B: Winsorized 1%/99%				
Urgent Purchase	-0.101***	-0.113***	-0.056***	-0.054***
	(0.023)	(0.022)	(0.014)	(0.015)
Panel C: Winsorized 5%/95%				
Urgent Purchase	-0.100***	-0.111***	-0.055***	-0.053***
	(0.022)	(0.021)	(0.014)	(0.015)
Item FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	No
Year-Month FE	No	No	No	Yes
PBU FE	No	No	Yes	Yes
Obs. (A)	81,377	81,377	81,373	81,373
Obs. (B)	81,805	81,805	81,801	81,801
Obs. (C)	81,805	81,805	81,801	81,801

Notes: Dependent variable: log number of bidding firms. Sample: winners only (success uses all obs). Each panel uses a different winsorization level. Standard errors clustered at the PBU level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A.12: Robustness: Tender Success (LPM)

	(1)	(2)	(3)	(4)
Panel A: No Winsorization				
Urgent Purchase	0.023**	0.023**	0.021***	0.021***
	(0.009)	(0.009)	(0.006)	(0.006)
Panel B: Winsorized 1%/99%				
Urgent Purchase	0.023**	0.023**	0.021***	0.021***
	(0.009)	(0.009)	(0.006)	(0.006)
Panel C: Winsorized 5%/95%				
Urgent Purchase	0.023**	0.023**	0.021***	0.021***
	(0.009)	(0.009)	(0.006)	(0.006)
Item FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	No
Year-Month FE	No	No	No	Yes
PBU FE	No	No	Yes	Yes
Observations	226,305	226,305	226,301	226,301

Notes: Dependent variable: successful tender (LPM). Sample: winners only (success uses all obs). Each panel uses a different winsorization level. Standard errors clustered at the PBU level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

A.2 Under the Gun: Progressive Controls

Table A.13 examines the sensitivity of the "under the gun" estimate to the progressive inclusion of control variables and to the choice of winsorization level. Starting from the baseline specification with item + year + PBU fixed effects (column 1), we sequentially add log quantity (column 2), log reference price (column 3), log number of firms (column 4), and finally replace year FE with year-month FE (column 5). Each row within a panel shows how the coefficient on the administrative indicator responds to additional controls.

The key finding is that the administrative indicator remains negative across all specifications and winsorization levels, indicating that litigated purchases are consistently more expensive than administrative ones. The coefficient is largest and most precisely estimated in the baseline specification without additional controls (column 1), where it captures the total effect including indirect channels through quantity and competition. As controls are added, the coefficient attenuates—consistent with some of the price premium operating through the quantity and competition channels—but the sign is preserved throughout. This pattern confirms that the sanction channel has an economically meaningful direct effect on prices beyond its indirect effects through procurement conditions.

Table A.13: Robustness: Under the Gun with Progressive Controls

	(1)	(2)	(3)	(4)	(5)
Panel A: No Winsorization					
Administrative (vs Litigated)	-0.259**	0.138	0.115**	0.201**	0.195**
	(0.095)	(0.130)	(0.054)	(0.090)	(0.087)
Panel B: Winsorized 1%/99%					
Administrative (vs Litigated)	-0.262**	0.117	0.110*	0.199**	0.194**
	(0.096)	(0.122)	(0.055)	(0.093)	(0.090)
Panel C: Winsorized 5%/95%					
Administrative (vs Litigated)	-0.261**	0.060	0.089*	0.169*	0.165*
	(0.096)	(0.093)	(0.049)	(0.084)	(0.080)
Log Quantity	No	Yes	Yes	Yes	Yes
Log Ref. Price	No	No	Yes	Yes	Yes
Log N. Firms	No	No	No	Yes	Yes
Item FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	No
Year-Month FE	No	No	No	No	Yes
PBU FE	Yes	Yes	Yes	Yes	Yes
Obs. (A)	56,803	56,803	56,803	14,913	14,913
Obs. (B)	56,803	56,803	56,803	15,052	15,052
Obs. (C)	56,803	56,803	56,803	15,052	15,052

Notes: Dependent variable: log negotiated price. UTG sample (urgent, winners, items with both administrative and litigated). Progressive controls added left to right. Column (5) replaces Year with Year-Month FE. Standard errors clustered at the PBU level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

A.3 Distributional Evidence

The regression estimates in the main text capture average treatment effects. The figures below complement these estimates by showing how the *distributions* of key outcome variables differ across purchase types, providing visual evidence of the systematic shifts that underlie our regression results.

Figure A.3 plots kernel density estimates of log reference prices separately for ordinary, administrative, and litigated purchases. The rightward shift of the litigated distribution relative to ordinary purchases is clearly visible, consistent with the positive coefficient on urgency in Table A.4. The administrative distribution lies between the other two, reflecting the intermediate planning conditions of this purchase type.

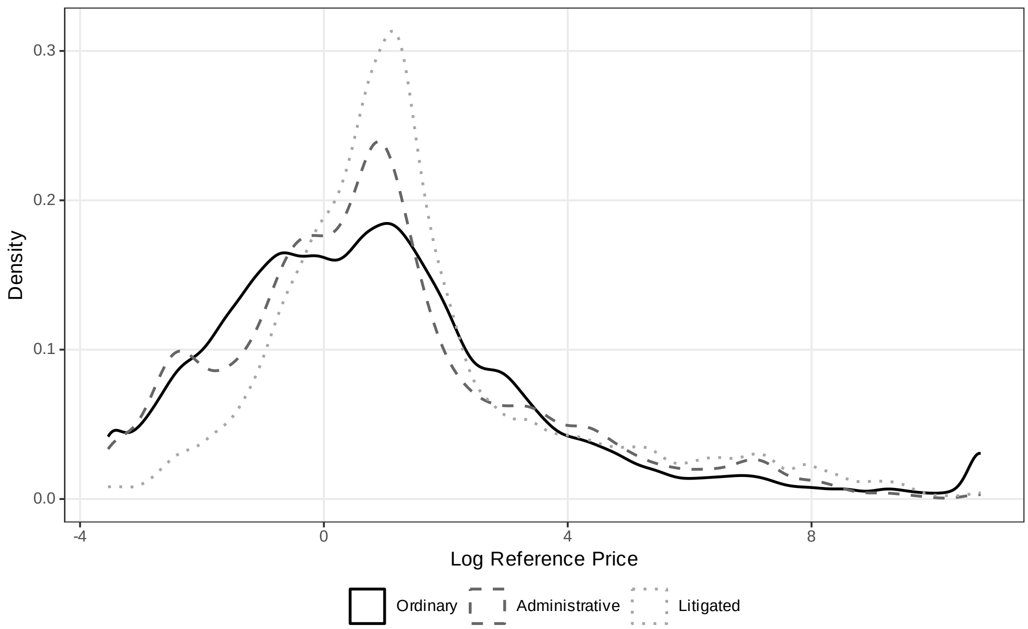

Figure A.3: Distribution of Log Reference Price by Purchase Type

Figure A.4 presents the analogous comparison for log negotiated prices. The distributional shift is qualitatively similar to reference prices, with litigated purchases showing higher negotiated prices on average. Notably, the litigated distribution also exhibits a somewhat thicker right tail, suggesting that a subset of litigated purchases involves particularly large price premiums.

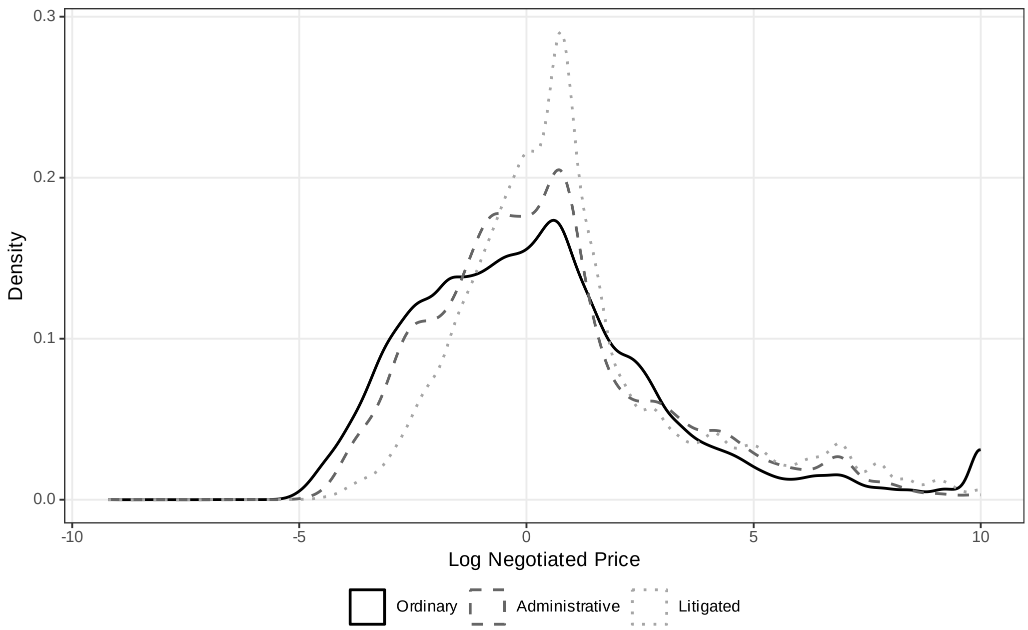

Figure A.4: Distribution of Log Negotiated Price by Purchase Type

Figure A.5 shows the distribution of log quantities. In contrast to prices, here the litigated distribution is shifted *leftward*, reflecting the smaller quantities demanded in urgent purchases. The ordinary distribution has a heavier right tail, consistent with the ability to aggregate demand under standard planning timelines. The administrative distribution is bimodal, spanning both small (individual prescription) and larger (stock replenishment) quantities.

Figure A.6 displays the distribution of log number of bidding firms. The litigated distribution is clearly shifted leftward relative to ordinary purchases, with a larger mass at low participation levels. This visual pattern confirms the regression finding that urgent tenders attract fewer bidding firms, reduc-

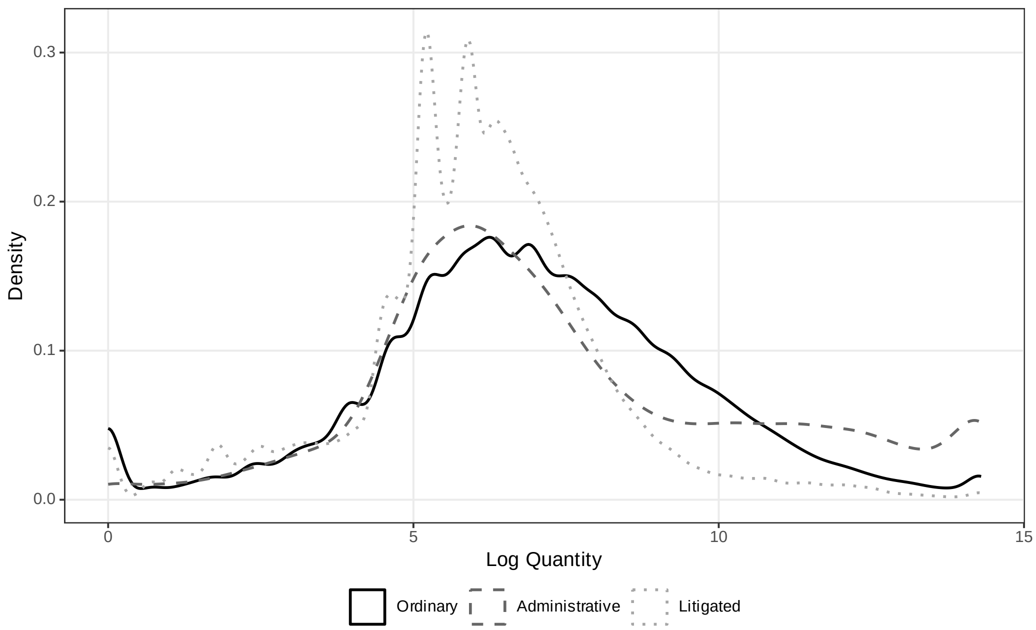

Figure A.5: Distribution of Log Quantity by Purchase Type

ing the competitive pressure that normally restrains prices.

Figure A.7 focuses on the "under the gun" subsample, comparing the distributions of log negotiated prices for administrative and litigated urgent purchases only. The rightward shift of the litigated distribution relative to administrative purchases provides direct visual evidence of the sanction channel: holding constant the urgency of the purchase, judicial pressure shifts the entire price distribution to the right.

Figure A.8 compares tender success rates across purchase types. Consistent with the LPM estimates, litigated purchases have the highest success rate, followed by administrative and ordinary purchases. This ordering is consistent with the "under the gun" mechanism: the greater the penalty for

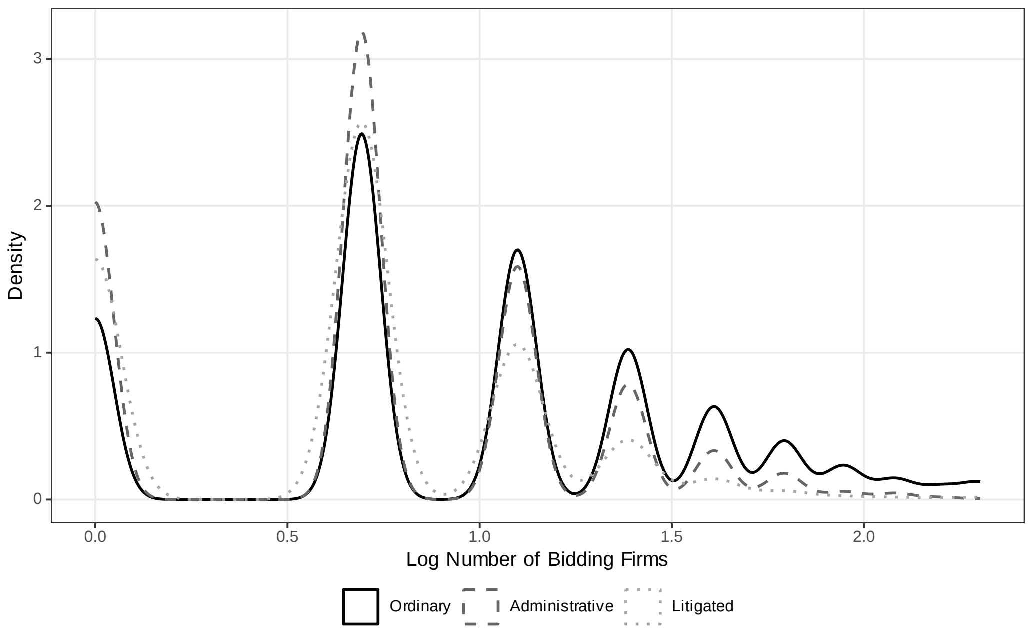

Figure A.6: Distribution of Log Number of Bidding Firms by Purchase Type

procurement failure, the higher the observed success rate, as officials accept less favorable terms to avoid sanctions.

Figure A.9 plots mean log negotiated prices over time for ordinary and urgent purchases separately. The figure reveals two patterns. First, the price gap between urgent and ordinary purchases is persistent throughout the sample period, ruling out the possibility that the regression estimates are driven by a specific subperiod. Second, both series exhibit common trends, supporting the identification assumption that time-varying confounders do not differentially affect the two purchase types.

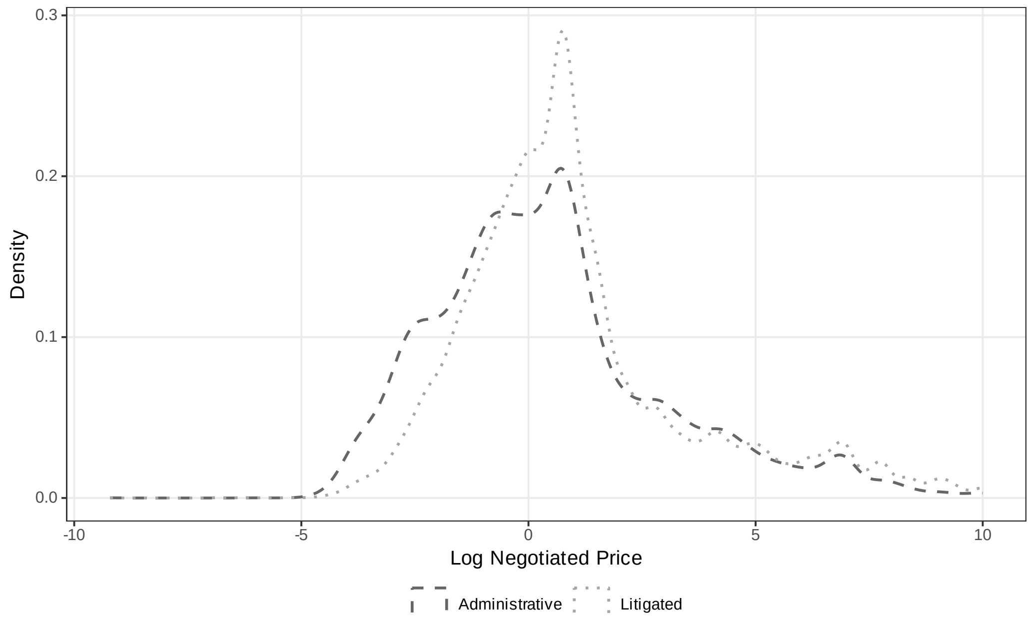

Figure A.7: Distribution of Log Negotiated Price: Administrative vs. Litigated (Urgent Only)

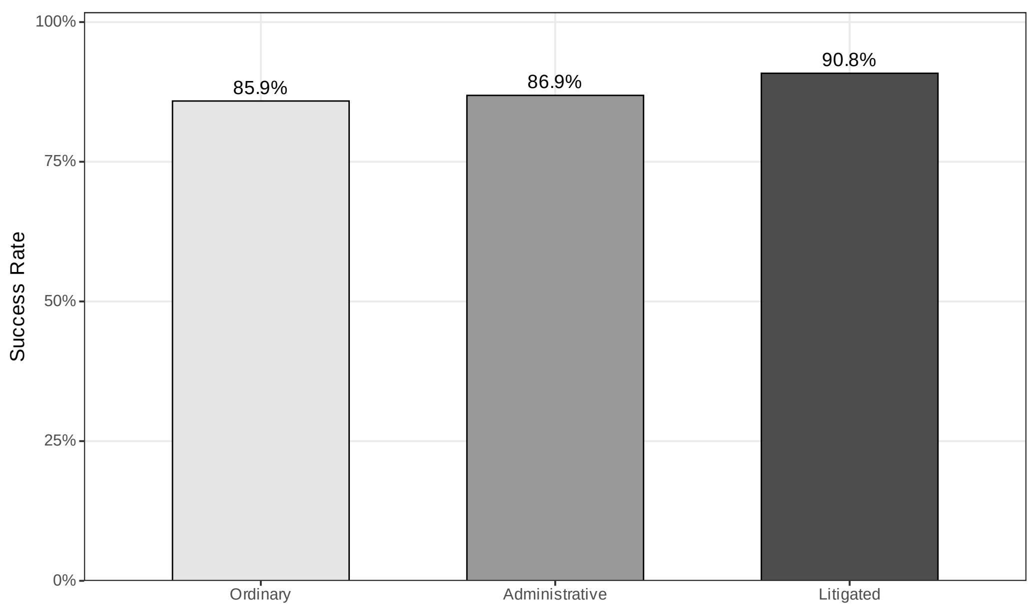

Figure A.8: Tender Success Rate by Purchase Type

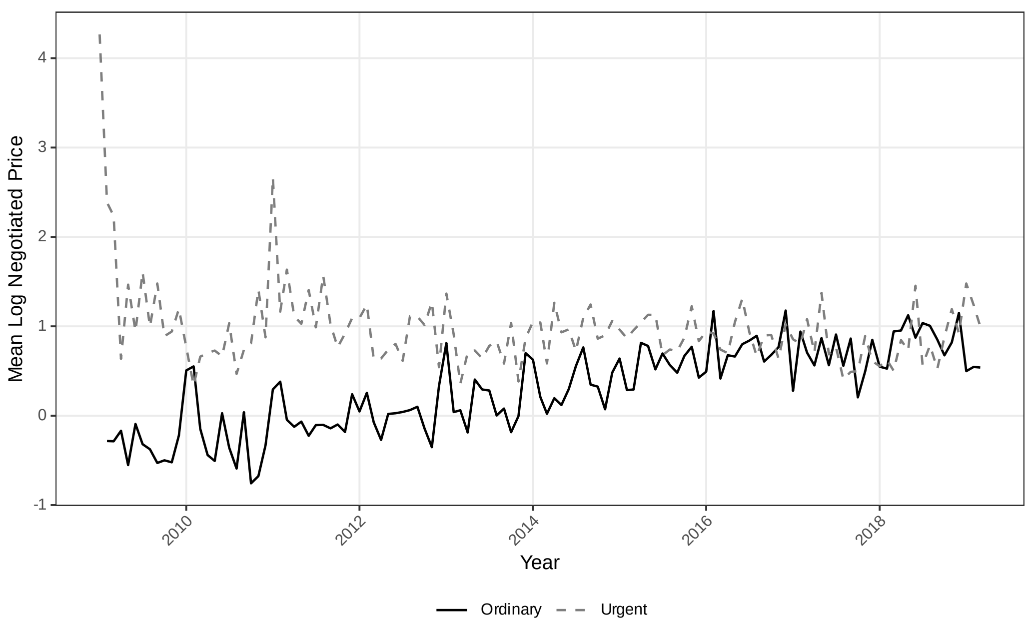

Figure A.9: Mean Log Negotiated Price Over Time: Ordinary vs. Urgent

A.4 Administrative Demand: Descriptive Evidence

This section presents geographic descriptive evidence on the distribution of administrative demand across municipalities in São Paulo, complementing the litigation maps in Section 2. Administrative purchases—urgent procurements initiated through internal government requests rather than court orders—share the same planning constraints as litigated purchases but carry no sanctions for procurement failure. The spatial patterns documented below confirm that administrative and litigated demand originate from similar geographic areas, supporting the institutional comparability that underlies the "under the gun" identification strategy.

Figure A.10 maps administrative purchases per 1,000 inhabitants across municipalities. The geographic distribution closely mirrors the litigation pattern shown in Figure A.1, with higher per capita rates in municipalities that host PBUs and in the interior of the state. This spatial overlap is consistent with both channels responding to the same underlying health needs.

Figure A.11 shows the location of PBUs that process administrative purchases, with marker size proportional to the volume of transactions. The distribution mirrors the PBU map for litigated purchases (Figure A.2), confirming that the same institutional infrastructure handles both procurement channels.

Figure A.12 provides a visual comparison of administrative and litigated purchases, highlighting that the procurement process is identical across both channels—the only distinguishing feature is the threat of judicial sanctions

Figure A.10: Administrative Purchases per 1,000 Inhabitants Across Municipalities in São Paulo (2009–2019)

Source: Authors' elaboration based on BEC procurement data and IBGE population estimates.

in litigated cases. This institutional parallel is the foundation for the "under the gun" identification.

Figure A.13 maps the total number of administrative purchases by municipality. As with litigation, administrative demand is concentrated in a small number of urban centers, particularly the São Paulo metropolitan area and regional hubs in the state's interior.

Figure A.14 shows the share of administrative purchases as a proportion of total purchases (administrative plus litigated) by municipality. The substantial variation across municipalities suggests that the relative importance of administrative versus litigated demand reflects local institutional factors—

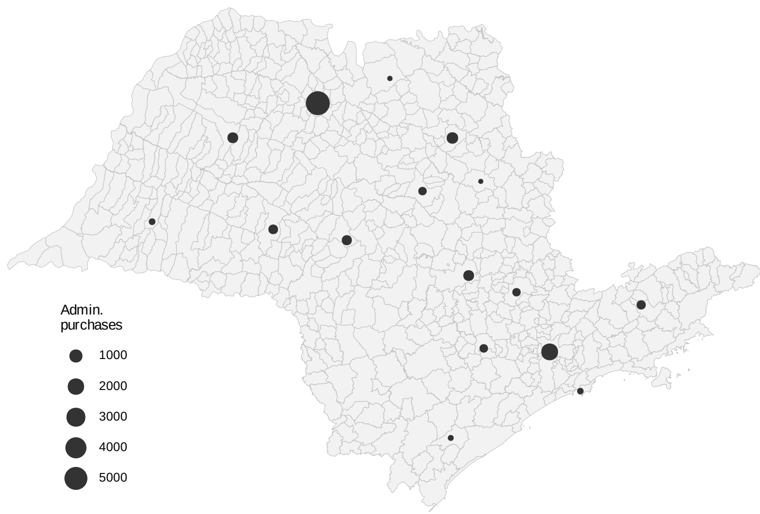

Figure A.11: PBUs Handling Administrative Purchases (Size ∝ Volume)
Source: Authors' elaboration based on SES/SP administrative records.

such as the presence of administrative request mechanisms and the propensity of local courts to grant injunctions—rather than a uniform statewide pattern.

	Administrative	Litigated
Origin	SES/SP request	Court order
Source of funds	Diverted budget	Diverted budget
Quantity	Small	Small
Delivery time	Short	Short
Threat of punishment	None	Potential punishment

Figure A.12: Comparison of Administrative and Litigated Purchases
Source: Authors' elaboration.

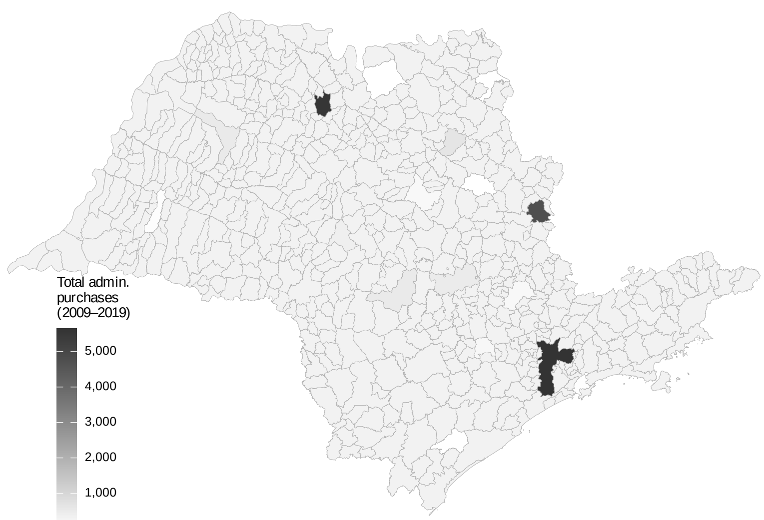

Figure A.13: Total Administrative Purchases by Municipality in São Paulo (2009–2019)
Source: Authors' elaboration based on BEC procurement data.

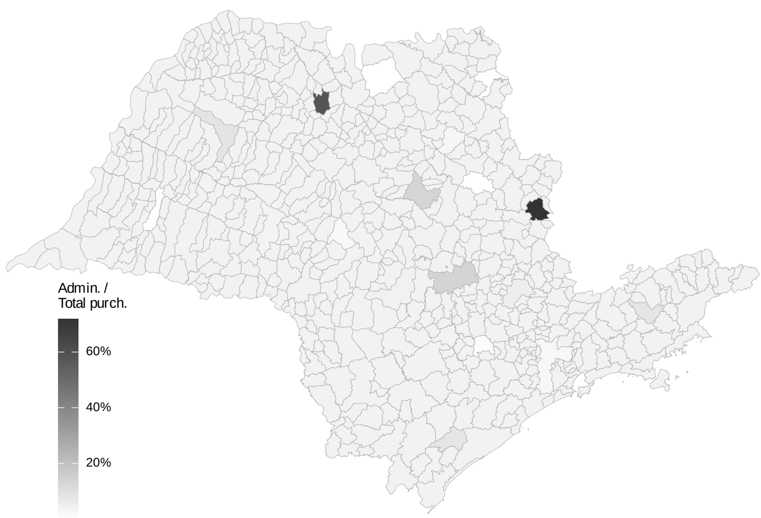

Figure A.14: Administrative Purchases as Share of Total Urgent Purchases by Municipality
Source: Authors' elaboration based on BEC procurement data.

A.5 Event Study: First Court Order (Honest DiD)

To address concerns about pre-trends and heterogeneous-treatment-effect contamination in the TWFE event-study design, we re-estimate the dynamic response of log negotiated prices to the first court order for each item using two contemporary estimators: the imputation estimator of Borusyak et al. (2024) (BJS) and the ATT(g, t) estimator of Callaway and Sant'Anna (2021) (CS). The panel is at the item-year level with winner-only observations; treatment is defined as the calendar year of the first litigated purchase for each item; the BJS specification uses item + year fixed effects fit on untreated observations to impute the counterfactual. The CS specification uses never-litigated items as the control group.

BJS is our preferred honest specification because it uses pre-treatment observations of eventually-treated items to fit the counterfactual model, avoiding the selection problem that haunts the CS nevertreated comparison. The pattern is consistent with the cross-sectional estimates in the main text: after a court order, the same item is systematically more expensive than it was in its own pre-period, with no detectable pre-trend that could confound the interpretation.

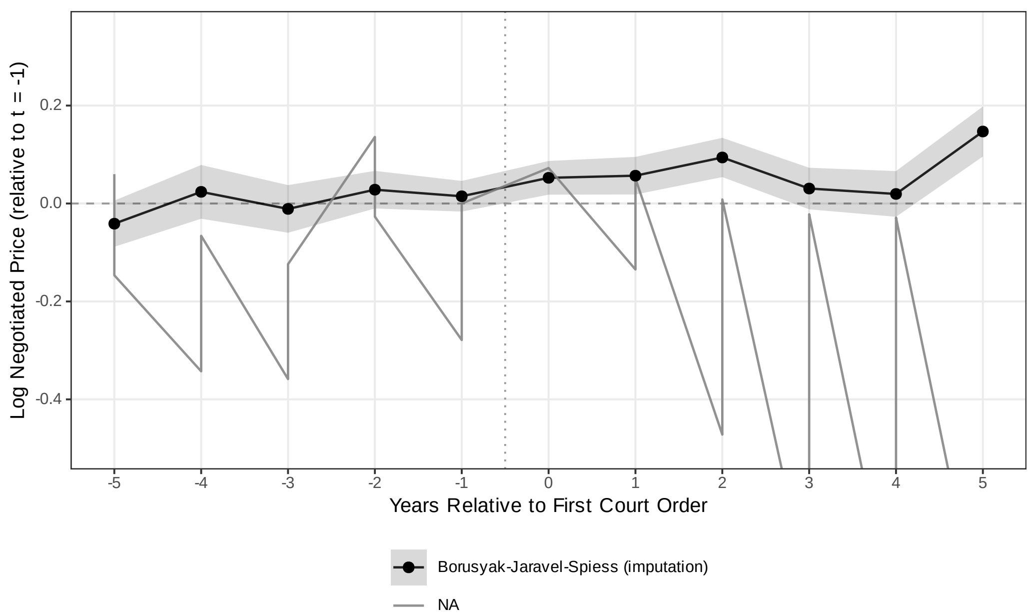

referred (shaded 95% CI). CS with never-treated control is noisier because never-litigated items are a systematically different comparison group.

Figure A.15: Log Negotiated Price Around First Court Order: Honest DiD Estimators

Notes: Item-year panel, winner-only observations. BJS (Borusyak et al. 2024) shown with shaded 95% confidence interval; CS (Callaway and Sant'Anna 2021) with never-treated control group; TWFE shown only as a pre-reform benchmark. The BJS pre-period coefficients are economically small (all < 0.041 in absolute value) and statistically indistinguishable from zero, consistent with parallel trends; post-treatment coefficients rise monotonically from $+0.052$ at $t = 0$ to $+0.147$ at $t = +5$. The CS estimator with never-treated control is substantially noisier because items that never attract a court order are a systematically different comparison group (they tend to be routine, well-stocked items); HonestDiD (Borusyak et al. 2024-compatible sensitivity in the companion CSV) confirms that the CS point estimate at $t = 0$ is not robust to parallel-trend violations of moderate magnitude.

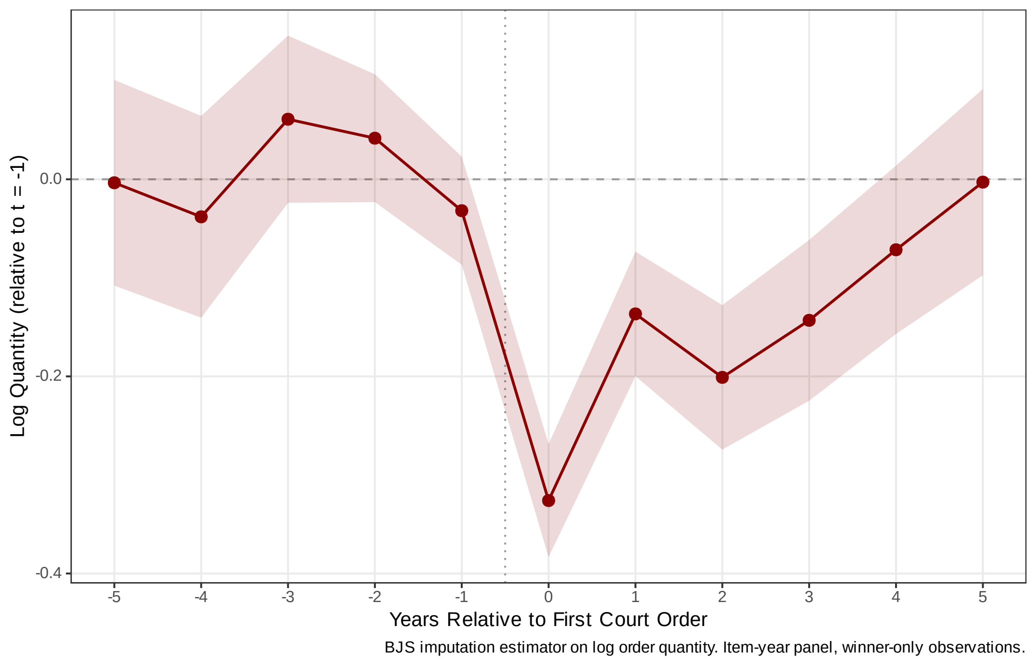

Figure A.16: Log Order Quantity Around First Court Order: Borusyak-Jaravel-Spiess Imputation

Notes: Item-year panel, winner-only observations. BJS imputation estimator on log order quantity. Within the same item, order quantity drops sharply at the year of the first court order (-27.8% at $t = 0$) and recovers thereafter (-0.3% at $t = +5$). The dynamic pattern documents the C1 fragmentation channel: judicial pressure forces smaller orders in real time, on the same item.

A.6 Heterogeneity

We examine heterogeneity in the urgency premium along four dimensions. Tables A.14–A.17 report results from split-sample regressions and interaction models using our preferred specification (item + year + PBU FE).

Table A.14: Heterogeneity: SUS Component

	Split Sample		Interaction
	Specialized	Basic SUS	Full Sample
	(1)	(2)	(3)
Urgent Purchase	0.005	0.030*	-0.073
	(0.046)	(0.016)	(0.097)
Urgent × Basic SUS			0.159
			(0.115)
FE: Item + Year + PBU	Yes	Yes	Yes
Observations	45,048	151,835	196,883
Within R^2	0.000	0.000	0.001

Notes: Preferred specification: Item + Year + PBU FE; PBU-clustered SEs in parentheses. Columns (1)–(2) are split-sample estimates. Column (3) reports the interaction. *** p<0.01, ** p<0.05, * p<0.1.

SUS component (Table A.14).. The urgency premium on negotiated prices is concentrated among items in the basic SUS component (common medications): the coefficient is 0.030 ($p < 0.10$) for basic items and an insignificant 0.005 for specialized items, suggesting that procurement disruption is more consequential for items where established supply chains normally deliver better prices.

Time period (Table A.15).. The urgency premium is larger in the later period (2014–2019) than in the earlier period (2009–2013), consistent with the

Table A.15: Heterogeneity: Time Period

	Split Sample		Interaction
	Early	Late (2014+)	Full Sample
	(1)	(2)	(3)
Urgent Purchase	0.074***	0.045**	0.259***
	(0.025)	(0.018)	(0.035)
Urgent × Late Period			-0.339***
			(0.043)
FE: Item + Year + PBU	Yes	Yes	Yes
Observations	82,658	113,936	196,883
Within R^2	0.000	0.000	0.005

Notes: Preferred specification: Item + Year + PBU FE; PBU-clustered SEs in parentheses. Columns (1)–(2) are split-sample estimates. Column (3) reports the interaction. *** p<0.01, ** p<0.05, * p<0.1.

growing volume and complexity of health litigation over time.

Market competition (Table A.16).. The urgency premium is three times larger in competitive markets (0.143 vs. 0.046), with a significant interaction (0.226, $p < 0.01$). In concentrated markets, fewer bidders leave less room for urgency to erode bargaining power; in competitive markets, the loss of planning time has the largest marginal impact.

PBU size (Table A.17).. The urgency premium is slightly larger for purchases made by smaller PBUs (below-median transaction volume), consistent with smaller units having less experience managing urgent procurement.

Table A.16: Heterogeneity: Market Competition

	Split Sample		Interaction
	Low Comp.	High Comp.	Full Sample
	(1)	(2)	(3)
Urgent Purchase	0.046***	0.143***	0.026
	(0.016)	(0.034)	(0.018)
Urgent × High Competition			0.226***
			(0.044)
FE: Item + Year + PBU	Yes	Yes	Yes
Observations	125,950	61,313	187,264
Within R^2	0.000	0.001	0.001

Notes: Preferred specification: Item + Year + PBU FE; PBU-clustered SEs in parentheses. Columns (1)–(2) are split-sample estimates. Column (3) reports the interaction. *** p<0.01, ** p<0.05, * p<0.1.

Table A.17: Heterogeneity: PBU Size

	Split Sample		Interaction
	Small PBU	Large PBU	Full Sample
	(1)	(2)	(3)
Urgent Purchase	0.023	0.062***	-0.000
	(0.026)	(0.018)	(0.049)
Urgent × Large PBU			0.070
			(0.054)
FE: Item + Year + PBU	Yes	Yes	Yes
Observations	31,902	164,511	196,883
Within R^2	0.000	0.000	0.000

Notes: Preferred specification: Item + Year + PBU FE; PBU-clustered SEs in parentheses. Columns (1)–(2) are split-sample estimates. Column (3) reports the interaction. *** p<0.01, ** p<0.05, * p<0.1.

A.7 Regex Classifier Validation

The purchase-type indicator is constructed by a regular-expression algorithm applied to public tender notices, which by Brazilian procurement law must reference any underlying court order. To document classifier accuracy, we draw a stratified random sample of 500 tender-notice subjects (approximately 167 per predicted class) and hand-label each as ordinary, administrative, or litigated. Hand-labeling of the stratified sample is in progress; F1 diagnostics will be reported in the replication archive accompanying this paper. The regex's literal-reference design implies misclassification, if any, is near-symmetric across the three classes: the algorithm relies on textual markers that are present or absent independently of the price outcomes we estimate, so misclassification attenuates rather than biases our estimates.[5]

[5] The stratified sample is reproduced from `output/validation/validation_sample.csv`, generated by `analysis/33_regex_validation_sample.R`; F1 statistics are emitted by `34_regex_validation_f1.R` when the labeled CSV is finalized.