Sourcing under Sanctions: Judicial Urgency and Pharmaceutical Procurement Costs*

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

Court mandates secure access to medicines, but they can also change how governments buy. We study São Paulo pharmaceutical procurement on BEC, covering 479,330 purchase-offer-item observations from 2009–2019. Higher prices under legal urgency can reflect two very different margins: incumbent suppliers charging more to sanctioned buyers, or fragmented sourcing that changes scale and supplier matching. Because court orders originate outside procurement offices, we interpret ordinary-versus-urgent estimates as the procurement effect of externally imposed legal urgency, conditional on item, time, and purchasing-unit fixed effects. A selected administrative urgent channel provides the closest feasible comparison without court sanctions; Lee bounds place the litigated-over-administrative price gap between 15.9% and 21.1%. The cost margin is not mainly a broad same-firm markup in deep repeated urgent markets: within the same firm, buyer, and item, prices are statistically indistinguishable across urgent regimes. Instead, judicial urgency operates through sourcing. Administrative orders are 3.3 times larger, and modal winners differ in 70.2% of item-buyer pairs. A residual within-firm price gap persists in thinner or earlier markets. The policy margin is not weaker access, but procurement capacity that preserves aggregation and supplier matching under legal urgency.

Keywords: public procurement, health litigation, accountability, sourcing

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

*Short paper. This version: May 2026. All errors are our own.

*Corresponding author.

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

1. Introduction

Courts can force governments to deliver public services. They can also change how the state produces those services. In right-to-health litigation, a judge orders a public agency to provide a specific medicine to an individual plaintiff. The agency must then source the medicine quickly, often under short deadlines and one-sided penalties for nondelivery. Court mandates do not merely change how much the state pays; they change how the state is forced to buy.

This production margin is distinct from the usual fiscal question in health litigation. A court order may have high private value and still be costly to execute because it arrives outside the normal procurement cycle. Conversely, a high observed price need not imply supplier exploitation: it may reflect smaller lots, compressed search, or a different supplier match. The empirical challenge is to distinguish the legal mandate from the procurement mechanism through which the mandate is fulfilled.

We study this question in São Paulo pharmaceutical procurement. The key institutional fact is simple: court orders originate outside the procurement process. Patients, physicians, lawyers, and judges determine whether an order arrives, for which medicine, and when. Purchasing units do not choose court-mandated status as a procurement strategy. This institutional sequence supports interpreting urgent purchases as externally imposed demand shocks to procurement, conditional on item, time, and purchasing-unit fixed effects. This is not a claim that patient or medicine selection into litigation is ignorable. It is a claim about the procurement transaction: the legal process moves demand into urgent execution from outside the purchasing unit.

The identifying variation is institutional rather than experimental: we compare procurement regimes for similar items, buyers, and periods after a legal process has shifted demand from ordinary planning into urgent execution.

The setting also offers a rare second comparison. Since 2009, São Paulo's health administration has operated an administrative request channel through which some patients obtain medicines outside court. These requests can generate urgent procurement without judicial sanctions. The channel is useful because it creates urgent pharmaceutical demand outside the court-sanction regime. It is limited because requests are screened by a scientific committee and are larger on average. We therefore treat administrative-versus-litigated procurement as a bounded comparison within urgent

procurement, not as unselected sanction exposure.

The empirical strategy follows the institution. Ordinary-versus-urgent comparisons estimate the procurement effect of externally imposed legal urgency. Administrative-versus-litigated comparisons ask how much of the within-urgent gap remains once administrative selection is bounded using Lee trimming. Mechanism tests then distinguish same-firm pricing from fragmented sourcing: quantity is treated as a post-treatment channel, within firm-buyer-item comparisons test whether incumbent suppliers charge more for the same item to the same buyer, and winner-switching evidence measures supplier-set reallocation.

The design is deliberately modest about what each comparison can carry. The broad urgent comparison has the strongest institutional claim. The administrative comparison is closer on the operational environment but weaker on selection. The within-firm comparison is sharp about pricing but silent about which suppliers win. Keeping these claims separate avoids turning the administrative channel or the mechanism regressions into more than they can support.

Four findings follow. First, urgent procurement is costlier and less competitive: relative to ordinary purchases of comparable items by the same buyers, urgent purchases have 5.4% higher negotiated prices and 5.4% fewer bidders, while tender success rises by 2.1 pp. Second, within urgent procurement, court-mandated purchases are more expensive than administrative urgent purchases. Because the administrative channel is selected, the preferred estimate is a Lee-bounded litigated-over-administrative gap of 15.9% to 21.1% rather than the naive gap of 29.5%. Third, there is little evidence of a broad same-firm markup in deep repeated urgent markets: within firm-buyer-item triples observed under both urgent regimes, the administrative coefficient is 0.035 (SE 0.041). Fourth, the sourcing margin is large. Administrative urgent orders are 3.3 times larger than litigated urgent orders; among item-buyer pairs observed under both regimes, modal winners differ in 70.2% of pairs and mean winner-set Jaccard similarity is 0.268.

The contribution is to open the procurement black box of judicial enforcement. Existing literatures study procurement waste, bureaucratic accountability, courts, and right-to-health litigation. This paper links them by separating two margins that standard procurement comparisons confound: same-firm pricing and supplier-set reallocation. In deep repeated urgent markets, court-mandated delivery raises costs through fragmented sourcing rather than broad same-firm markups. This reframes right-to-health litigation as a public-sector production problem under legal constraint,

connecting procurement research on price dispersion, buyer capacity, and passive waste (Bandiera et al., 2009; Best et al., 2023; Bosio et al., 2022; Decarolis et al., 2020), courts and procurement performance (Coviello et al., 2018), and accountability, bureaucratic distortion, and health litigation (Prendergast, 2007; Bandiera et al., 2021; Ferraz, 2009; Wang, 2015; Biehl et al., 2009). The policy implication is not to weaken access to medicines. It is to preserve delivery while rebuilding aggregation under legal urgency.

2. Institutional Background

2.1. Judicial orders as procurement shocks

Brazil's constitutional right to health gives courts a central role in disputes over access to medicines (Ferraz, 2009; Wang, 2015; CNJ/INSPER, 2019). In São Paulo, patients can obtain court orders requiring the state health administration to provide a specific medicine. The legal process is initiated outside the procurement office: patients, physicians, lawyers, public defenders, and judges determine whether an order arrives, what medicine it covers, and when the agency must respond. The institutional fact that matters for procurement is not the legal doctrine itself, but the hard delivery obligation it creates. The purchasing unit receives the mandate as an obligation to source, not as a procurement option it selected.

This distinction is operational as well as legal. A procurement unit can choose how to write a tender, when to aggregate ordinary demand, and which procurement modality to use within statutory rules. It cannot choose whether a plaintiff files a case, whether a judge grants relief, or whether the order names a medicine that must be delivered on a short schedule. Once the order arrives, the unit's choices are downstream: how to comply, which suppliers to attract, and whether to accept a costly outcome to avoid nondelivery.

This institutional sequence disciplines the paper's first comparison. Court-mandated status is imposed on public buyers, not chosen to obtain a favorable price, select a supplier, or improve a tender outcome. The remaining concern is that the health-demand process selects different patients or medicines into litigation. The design addresses that concern with item, time, and buyer controls, administrative-versus-litigated bounds, and mechanism evidence rather than by claiming that selection into litigation is ignorable.

Judicial orders also change the procurement environment. They often arrive with short deadlines and compliance measures, including daily fines or fund seizure for nondelivery. The resulting pressure is one-sided: the public buyer must deliver, while suppliers remain voluntary participants who choose whether to bid and at what price. Procurement notices may reference the court order or associated legal process, so bidders can sometimes infer that the buyer is under judicial urgency.

2.2. Procurement on BEC

São Paulo's Department of Health buys medicines through BEC, the state electronic procurement platform. BEC records item identifiers, buyers, quantities, reference prices, negotiated prices, procurement modalities, participating firms, winning suppliers, and tender outcomes. Ordinary purchases, administrative urgent purchases, and litigated urgent purchases are therefore observed within the same procurement infrastructure.

Ordinary procurement can plan recurring demand, consolidate orders, and wait for broader supplier participation. Urgent procurement compresses this process. A buyer responding to a patient-specific mandate may buy a smaller lot, repeat tenders more frequently, or accept a thinner supplier set. Judicial urgency therefore has direct sourcing consequences. Because BEC observes all three regimes on a common platform, the data can separate same-firm pricing, quantities, and winner switching.

2.3. Administrative urgent demand

The administrative-request channel processes some patient demands outside court. Patients can request medicines directly through the health administration; accepted requests can generate urgent individualized procurement without a judicial order. These purchases share the operational pressure of individualized delivery but do not carry the court-sanction regime described above.

The channel is valuable precisely because it is imperfect. It creates urgent pharmaceutical demand outside court sanctions, but accepted requests pass through scientific and cost-effectiveness screening and can differ from litigated cases in item mix, feasibility, and scale. The paper therefore uses administrative purchases as the closest feasible urgent-procurement comparison, not as a clean counterfactual. Ordinary purchases reflect planned demand; administrative purchases reflect screened urgent requests without court sanctions; litigated purchases reflect court-mandated urgent delivery under possible fines, seizure, or liability.

This conservative use of the administrative channel is central. A direct comparison would attribute the entire litigated-administrative gap to sanctions, but the committee can admit cases that are more feasible or more cost-effective. We instead ask whether a positive gap remains after allowing for monotone selection into the administrative channel, and through which procurement margin it operates.

3. Data, Classification, and Samples

3.1. Procurement and litigation data

The data cover São Paulo pharmaceutical procurement through BEC during 2009–2019. The empirical file contains 479,330 purchase-offer-item observations. BEC records item identifiers, buyers, suppliers, reference prices, negotiated prices, quantities, participating firms, winners, procurement modality, and tender outcomes. We link these procurement records to SES/SP litigation and administrative-process markers to classify purchases into ordinary, administrative urgent, and litigated urgent regimes.

The classifier operates upstream at the purchase-order/tender-notice level, where legal status is observed. After regime labels are linked to BEC item records, outcomes are measured at the purchase-offer-item level. Price regressions use accepted winning bids; other outcomes use the broader purchase-offer-item or tender-level sample specified in each table. This preserves the level of the legal mandate while using the transaction level at which prices, quantities, and winners are realized.

3.2. Classification of purchase regimes

Purchase regimes are classified using structured links and tender-notice text. The classifier combines machine-learning predictions, JSON fields, regex markers, and position-based checks to separate ordinary, administrative urgent, and litigated urgent purchases.

Validation is conducted at the purchase-order/tender-notice level. Against 179,148 ground-truth purchase orders, the classifier reaches 98.6% exact agreement. Per-class F1 is 0.93 for judicial purchases and 0.96 for administrative purchases; the macro-F1 over those two urgent classes is 0.94. Online Appendix A reports the validation and sample-construction tables.

Remaining misclassification would attenuate regime contrasts unless it is systematically correlated with residual prices within controls. The main concern is differential error between urgent regimes among high-cost items; the validation evidence reduces, but cannot eliminate, that concern.

3.3. Samples and identifying variation

The paper uses several samples because each empirical comparison asks a different question. The analysis sample supports ordinary-versus-urgent comparisons. The winners-only price sample contains 196,883 accepted winning bids and is used when negotiated price is the outcome.

Within urgent procurement, the urgent panel contains administrative and litigated winning bids and supports the under-the-gun comparison. The firm-buyer-item triple sample is narrower by design. It requires the same supplier to sell the same item to the same buyer under both urgent regimes, so it identifies within-supplier pricing. It does not measure supplier-set reallocation, which requires the winner-switching evidence in Section 5.

The samples narrow as the question narrows. The ordinary-versus-urgent comparison supports the broad urgent-demand claim. The urgent panel focuses on sanction exposure among purchases already under urgent execution. The triple sample conditions on the same supplier; that narrowness is the source of the pricing test, not an estimate of total costs.

3.4. Descriptive facts

The descriptive facts show why raw administrative-versus-litigated means are not a design. Administrative urgent purchases are larger than litigated urgent purchases, while litigated purchases have higher average log negotiated prices. These differences combine selection, scale, and sanction exposure, so we do not interpret raw means causally. The design instead bounds selection, decomposes scale, and tests same-firm pricing and supplier-set reallocation directly.

The descriptive patterns nevertheless motivate the mechanism analysis: larger administrative orders suggest more aggregation within the urgent tier, while higher raw litigated prices show the direction of the problem but not its cause.

4. Empirical Strategy

4.1. Hierarchy of estimands

The empirical strategy separates five objects. The first is the procurement effect of legal urgency externally imposed on the purchasing unit:

$$y_{oibt} = \beta \, \text{Urgent}_{oibt} + \alpha_i + \gamma_t + \delta_b + \varepsilon_{oibt}, \tag{1}$$

where o indexes the purchase-offer-item observation, i the item, b the buyer or purchasing unit, and t the year or month. The dependent variable is log negotiated price, log reference price, log number of bidders, or tender success. Price regressions use accepted winning bids; success regressions use the broader tender-outcome sample. Fixed effects absorb item, time, and buyer differences.

We interpret β as the procurement effect of legal urgency, conditional on item, time, and purchasing-unit fixed effects, under the identifying assumption that residual procurement shocks do not drive the timing and content of court orders within those cells. The institutional support is that court orders shift demand into urgent execution from outside the purchasing unit: procurement offices do not choose court-mandated status as a buying strategy.

The main threat is that unobserved service failures, stockouts, or local administrative breakdowns may both trigger litigation and raise procurement costs. The fixed effects map directly to this concern. Item effects absorb persistent differences across medicines, time effects absorb aggregate shocks to pharmaceutical markets and procurement, and buyer effects absorb stable differences in purchasing capacity. The comparison is therefore between different procurement regimes for comparable items bought by comparable buyers in comparable periods. The design also uses a placebo on never-litigated items, the selected administrative urgent comparison with Lee bounds, and mechanism evidence to check whether the price pattern is consistent with externally imposed legal urgency rather than generic procurement deterioration.

4.2. Sanction exposure and mechanisms

The second object is the within-urgent sanction margin. We compare administrative and litigated urgent purchases:

$$y_{oibt} = \theta \, \text{Admin}_{oibt} + \alpha_i + \gamma_t + \delta_b + \varepsilon_{oibt}, \tag{2}$$

where Admin_{oibt} equals one for administrative urgent purchases and zero for litigated urgent purchases. Coefficients are administrative minus litigated. In price regressions, a negative θ means litigated purchases are more expensive; tables also report percentage gaps in the reader-facing litigated-over- administrative direction.

Administrative purchases remove judicial sanctions but introduce selection. The scientific committee admits cases using medical and cost-effectiveness criteria, so administrative purchases can be easier to source, more standardized, or larger even absent sanctions. We therefore do not use θ as a point-identified sanction effect. We report a selection-bounded sanction margin using Lee trimming (Lee, 2009) within item $\times$ year $\times$ PBU strata. The trimming removes the overrepresented administrative observations from the high and low tails of the within-stratum price distribution, producing lower and upper bounds for the litigated-over-administrative price gap under monotone selection. This interval is the preferred estimate of the sanction margin.

The Lee exercise is deliberately conservative. It does not make the administrative channel experimental; it turns an implausible point estimate into a bounded rather than point-identified comparison with a transparent monotonicity assumption. The naive coefficient is reported only as a benchmark.

The third object is same-firm pricing:

$$y_{ofibt} = \kappa \, \text{Admin}_{ofibt} + \mu_{fib} + \gamma_t + \varepsilon_{ofibt}, \tag{3}$$

where f indexes the supplier and μ_{fib} is a firm-by-item-by- buyer fixed effect. This comparison uses observations where the same firm sells the same item to the same buyer under both urgent regimes. It tests same-firm pricing and intentionally conditions away supplier-set reallocation. We therefore pair it with direct sourcing evidence: order-size gaps, the decomposition figure, and winner-set overlap across regimes. The combination distinguishes incumbent pricing from fragmented sourcing.

The fourth object is sourcing mechanism evidence. Observed price gaps can arise because urgent purchases lose scale or because the winning supplier set changes. Quantity is post-treatment: urgency can change order size, and order size then affects prices through bulk discounts. Specifications that condition on log quantity are channel diagnostics, not preferred total-effect estimates. Winner-switching comparisons document supplier-set reallocation but do not alone price that re-

allocation. The decomposition is an accounting and mechanism framework, not a structural social valuation.

The fifth object is timing. Event-study evidence around first item exposure is used as a diagnostic for whether price changes line up with litigation exposure. It is not the primary identification design. The mechanism claim rests on the joint pattern: limited same-firm pricing evidence in deep markets, large quantity differences, and direct sourcing evidence.

5. Results

5.1. Judicial urgency raises procurement costs

Panel A of Table 1 compares urgent purchases with ordinary purchases of the same items by the same buyers over time. Negotiated prices are 5.4% higher, reference prices are 2.7% higher, and the number of bidding firms falls by 5.4%.[1] Tender success rises by 2.1 pp. The pattern is the procurement effect of externally imposed legal urgency under the identifying assumption stated above: the state completes more tenders, but under weaker competitive conditions and at higher prices.

The success result is useful for interpretation. Higher prices could arise mechanically if urgent tenders failed more often and only the most attractive urgent purchases survived into the price sample. Instead, urgent procurement succeeds more often despite fewer bidders and higher prices. That pattern is consistent with a buyer under delivery pressure accepting less favorable terms to complete the purchase. It also clarifies the role of the ordinary-versus-urgent comparison: the estimate is not simply a price premium among successful tenders, but part of a broader change in the procurement frontier. These estimates establish that legal urgency changes procurement, but they do not yet distinguish pricing from sourcing.

[1]The price premium and the reduction in bidders coincide at 5.4% only after rounding; the underlying log coefficients are 0.053 (price) and -0.056 (bidders).

Table 1: Urgent procurement margins and under-the-gun bounds.

Panel A. Urgent versus ordinary procurement

Outcome	Effect	SE	Interpretation
Log negotiated price	0.053	0.016	5.4% higher prices
Log reference price	0.027	0.014	2.7% higher reference prices
Log number of bidding firms	-0.056	0.014	5.4% fewer bidders
Tender success	0.021	0.006	2.1 pp higher success

Panel B. Under-the-gun: administrative vs. litigated urgent gap

Specification	Admin coef.	SE	Gap (%)	N
Naive UTG: item + year + PBU FE	-0.259	0.092	29.5%	61,620
Lee lower bound: admin bottom-tail trim	-0.148	0.097	15.9%	45,624
Lee upper bound: admin top-tail trim	-0.192	0.095	21.1%	45,624

Notes: POI denotes purchase-offer-item. *Panel A* reports urgent-versus-ordinary estimates under item, year, and PBU fixed effects; negotiated-price and reference-price estimates use accepted winning bids, tender success uses the broader tender/POI sample, and standard errors are clustered by PBU. Panel A establishes the urgent-procurement margin and is not the sanction-exposure design. *Panel B* reports the administrative-versus-litigated urgent gap; coefficients are administrative minus litigated log negotiated price, so negative coefficients mean litigated purchases are more expensive, while the percentage column is the reader-facing litigated-over-administrative gap. Lee trimming is applied within item×year×PBU strata where administrative observations exceed litigated observations; the mean trimming rate is 26.9% and the maximum is 100.0%.

5.2. The sanction margin is selection-bounded

Panel B moves inside urgent procurement. The naive litigated-over-administrative gap is 29.5%, but the administrative channel is screened and larger. After trimming the selected administrative group within item-by-year-by-PBU strata, the litigated-over-administrative price gap lies between 15.9% and 21.1%. The interval remains positive, so the court- sanction regime is associated with higher urgent procurement costs even after bounding the most direct selection concern. The estimate is not a structural counterfactual with sanctions removed; it is a selection-bounded sanction margin.

This result should be read together with Panel A. The first panel reports the procurement effect of legal urgency under the identifying assumption in Section 4. The second panel is a selected but informative comparison within urgent procurement: it is bounded rather than point-identified and is interpreted as a Lee-bounded sanction margin. The next question is whether that margin is same-firm pricing.

5.3. Same-firm pricing is not the broad margin in deep markets

Table 2 asks whether the same supplier charges more when selling the same item to the same buyer under the litigated regime. In the firm-buyer-item triple sample, the administrative coefficient is 0.035 with standard error 0.041. Because the coefficient is administrative minus litigated, the same-firm pricing test provides little evidence of a broad same-firm markup in this repeated urgent-market sample.

Table 2: Within firm-buyer-item null: robustness across subsamples.

Subsample	$\hat{\beta}_{\text{Admin}}$	SE	N
All triples (baseline)	0.035	0.041	4,573
Above-median quantity	-0.005	0.044	2,188
Below-median quantity	0.066***	0.025	2,027
SUS-formulary items	-0.001	0.026	2,540
Non-formulary items	0.101	0.064	2,033
Earlier period	0.117***	0.037	2,553
Later period	-0.038	0.052	1,492

Notes: Each row reports the within firm-buyer-item triple Admin coefficient on log negotiated price, with FBI and year fixed effects, PBU-clustered SEs. Coefficients are administrative minus litigated; negative values mean litigated purchases are more expensive within the same firm, buyer, and item, while values near zero indicate no detectable within-triple price difference. Subsamples are constructed from the triple sample (1,206 triples). Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

This is a demanding mechanism test. The same firm observes the same buyer and sells the same item under both urgent regimes. The near-zero estimate does not say that suppliers never exploit urgency. It says that the broad procurement-cost margin is not well explained by the same incumbent supplier charging much more for the same item to the same buyer in deep repeated urgent markets.

The null is concentrated in deeper repeated urgent markets. The coefficient remains near zero for above-median-quantity purchases and SUS-formulary items, where scale and standardization make repeated sourcing more likely. It becomes positive in below-median-quantity purchases and in the earlier period. The quantity split should not be read as same-firm pricing alone: the within-triple log-quantity coefficient is -0.259 (SE 0.074), so quantity captures the bulk-discount channel. The earlier-period gap is a genuine residual within-firm difference, but it does not point to litigated-buyer leverage: the surviving coefficient is administrative-dearer, not litigated-dearer, and it fades

as urgent markets mature.

The heterogeneity is economically sensible. Thin markets leave the buyer with fewer outside options and make short-deadline delivery harder to substitute across firms. Earlier-period procurement may also reflect a less mature administrative response to judicial demand. It locates where same-firm pricing could in principle matter. But the surviving within-firm gap is administrative-dearer and time-declining, so it reads as early-market maturation rather than court-order leverage over the sanctioned buyer. The absence of broad same-firm markups in deep markets shifts attention to sourcing.

5.4. The main mechanism is fragmented sourcing

The remaining evidence points to sourcing. Administrative urgent orders are 3.3 times larger than litigated urgent orders, and the bulk-discount elasticity is -0.329. Court-mandated buying therefore gives up scale. Figure 1 reconciles the observed administrative-minus-litigated price gap with a mechanical quantity component, the within-firm pricing component, and a residual composition component.

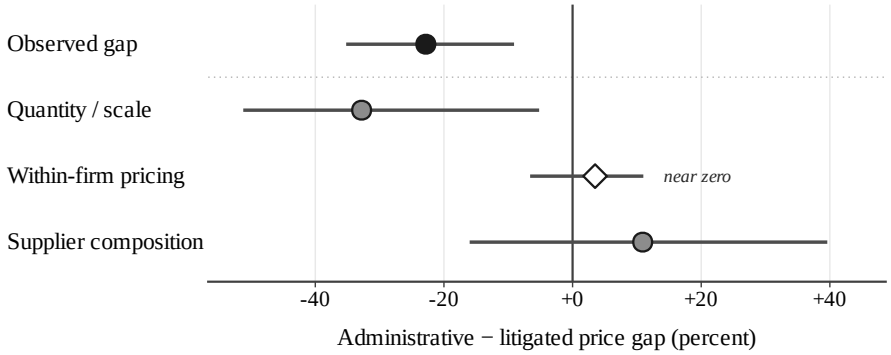

Figure 1: Pricing versus sourcing decomposition. Components are administrative minus litigated. The figure separates the observed price gap from the mechanical quantity component, the within firm-buyer-item pricing component, and a residual composition component. The residual is interpreted together with the winner-switching evidence in Table 3.

Table 3 provides the direct sourcing test. Among item-buyer pairs observed under both urgent regimes, mean winner-set Jaccard similarity is 0.268. 48.5% of pairs have no winner overlap, and the modal winner differs in 70.2% of pairs. Conditional on the same firm, prices are not broadly higher in deep markets; unconditionally, the state often buys from a different supplier set. This combination is the empirical signature of fragmented sourcing: near-zero within-firm pricing in deep markets and large supplier-set reallocation.

Table 3: Winner switching across urgent regimes within item-buyer pairs.

Item-buyer pairs (both regimes)	2,134
Mean distinct winners, admin	2.05
Mean distinct winners, litigated	1.99
Mean Jaccard similarity, winner sets	0.268
Pairs with any winner overlap (%)	51.5
Pairs with NO winner overlap (%)	48.5
Pairs with same modal winner (%)	29.8
Pairs with different modal winner (%)	70.2

Notes: Unit is a buyer×item pair with at least one administrative and one litigated urgent purchase. Jaccard similarity is the ratio of the cardinality of the intersection to the cardinality of the union of the two regimes' winning-firm sets. "Modal winner" is the most frequent winning firm within each regime for a given pair. Source: BEC-SP, 2009–2019.

The decomposition and winner-switching evidence explain why a policy focused only on posted prices would be incomplete. The state is not just buying the same thing from the same firm at a higher price; it is often buying a smaller lot from a different set of firms.

6. Falsification and Robustness

6.1. Placebo outside litigated items

The placebo uses items never observed in litigation. If the urgent- price pattern reflected generic BEC dynamics or broad time trends, administrative urgency among never-litigated items should reproduce it. It does not: the negotiated-price placebo coefficient is -0.020 with standard error 0.032. This supports the interpretation that the main urgent premium is tied to the litigation margin rather than to platform-wide procurement dynamics.

6.2. Inference and dynamics

Because PBU clusters are few and uneven, the appendix reports Rademacher wild-cluster inference for the under-the-gun contrast. The preferred specification rejects a zero gap ($p = 0.0080$), and the tighter item-by-year-month specification gives $p = 0.0390$. Dynamic evidence is used as a timing check: the BJS event-study estimates (Borusyak et al., 2024) rise after item exposure, but Honest-DiD sensitivity (Rambachan and Roth, 2023) is not strong enough to make the dynamic

design the primary identification strategy. The main claims rest on the externally imposed urgency design, the selection bounds, and the mechanism tests.

7. Policy Interpretation and Limits

7.1. What the evidence shows

The evidence identifies a procurement mechanism behind judicial enforcement. Court mandates secure delivery, but externally imposed urgency raises costs and weakens competition. Within urgent procurement, the sanction-related gap remains positive after selection bounds. In deep repeated urgent markets, however, the broad margin is not the same firm charging much more for the same item to the same buyer. It is fragmented sourcing: smaller orders, lost scale, and different winning suppliers.

The mechanism is a judicial-enforcement form of passive waste rather than a claim of corruption or active overpayment. One-sided sanctions help secure delivery, but they move the state into a worse sourcing technology. Procurement officials may be doing exactly what the court requires while buying under conditions that make success more expensive.

This distinction is important in a setting where the legal and fiscal stakes are both real. The estimates do not imply that officials should resist court orders, or that plaintiffs should be denied access. They show that a procurement system designed around planned, recurring demand performs differently when repeated patient-specific mandates arrive as urgent obligations. The relevant administrative problem is therefore not simply compliance. It is compliance without destroying the scale and supplier relationships that make public buying cheaper.

7.2. What the evidence does not show

The paper does not estimate patient health benefits or the full social value of right-to-health litigation. It also does not identify a structural counterfactual in which court sanctions are removed from otherwise identical litigated demand. The fiscal calculation in the appendix is a procurement-cost calculation only.

Nor does the paper rank legal enforcement against alternative health policies. A court order can be socially valuable even if it raises procurement costs; conversely, a low-cost procurement response says nothing about whether the underlying allocation of medicines is equitable. The estimates

isolate one margin for that broader evaluation: the extra procurement cost created when urgent legal demand disrupts aggregation and supplier matching.

7.3. Restoring aggregation under judicial urgency

The policy margin is to preserve access while preventing legal urgency from becoming procurement fragmentation. Framework agreements, pooled urgent procurement, pre-contracted suppliers, inventories for recurrent litigated medicines, and administrative screening capacity all target the measured mechanism: aggregation and supplier matching under urgent legal constraints, consistent with evidence on group purchasing in health-sector procurement (Lin and Wang, 2025). The paper does not estimate these reforms directly, but the mechanism evidence identifies why they are the relevant policy class.

8. Conclusion

Court orders are a powerful instrument for securing access to medicines. They also arrive as procurement shocks. In São Paulo, those shocks raise prices and reduce bidder participation, even as the state completes urgent tenders more often. The administrative urgent channel shows that the court-sanction margin remains positive after selection bounds, but the mechanism is not a broad incumbent markup in deep repeated urgent markets. The main loss is sourcing efficiency: judicial urgency fragments demand, reduces scale, and changes which suppliers win.

That distinction matters for policy. If the procurement-cost margin were mainly same-firm pricing, price caps or contract terms would be the natural tools. If the margin is fragmented sourcing, the relevant tools are those that preserve aggregation while respecting judicial delivery obligations: framework agreements, pooled urgent procurement, pre-contracted suppliers, inventories for recurrent medicines, and administrative screening capacity. The point is not to weaken access, but to prevent legal urgency from forcing the state to buy one patient-specific emergency at a time.

The estimates are specific to São Paulo pharmaceutical procurement during 2009–2019 and do not value health benefits or the full social value of litigation. They identify a narrower but policy-relevant margin: how judicial enforcement changes the way the state buys. For public economics, that margin is central. Legal rights are implemented through administrative production systems.

When enforcement changes those systems' timing, scale, and supplier matches, the fiscal cost of a right depends not only on what is mandated, but on how the state must procure it.

References

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., 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.

Bosio, E., Djankov, S., Glaeser, E.L., Shleifer, A., 2022. Public procurement in law and practice. American Economic Review 112, 1091–1117. doi:10.1257/aer.20200738.

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

Coviello, D., Moretti, L., Spagnolo, G., Valbonesi, P., 2018. 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.

Lee, D.S., 2009. Training, wages, and sample selection: Estimating sharp bounds on treatment effects. Review of Economic Studies 76, 1071–1102. doi:10.1111/j.1467-937X.2009.00536.x.

Lin, H., Wang, Y., 2025. The value of group purchasing: Evidence from the U.S. hospital industry. Journal of Public Economics 247, 105380. doi:10.1016/j.jpubeco.2025.105380.

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

Rambachan, A., Roth, J., 2023. A more credible approach to parallel trends. Review of Economic Studies 90, 2555–2591. doi:10.1093/restud/rdad018.

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.