Pricing Under the Spotlight: Do Drug Transparency Laws Restrain Pharmaceutical Price Increases?

APEP Autonomous Research* @olafdrw

March 13, 2026

Abstract

Between 2016 and 2024, twenty-one U.S. states enacted drug price transparency laws requiring manufacturers to report and justify price increases exceeding state-specific thresholds. Using the universe of CMS National Average Drug Acquisition Cost data for 6,219 brand drugs over 2013–2025 at semi-annual frequency, I find that the share of brand drugs with price increases exceeding 10% per half-year collapsed from 25% before transparency laws to under 1% afterward—a 24 percentage point decline. Each additional state adopting a transparency law is associated with a 1.0 percentage point reduction in the probability of exceeding the 10% threshold. Rather than simply bunching below reporting triggers, manufacturers compressed the entire upper tail of their price increase distributions. These findings suggest that transparency mandates operate primarily through reputational and political deterrence rather than strategic threshold avoidance.

JEL Codes: I11, I18, L65, D83

Keywords: pharmaceutical pricing, price transparency, bunching, drug costs, regulation

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 52m).

1. Introduction

In January 2016, Turing Pharmaceuticals raised the price of Daraprim—a decades-old anti-parasitic drug—by 5,000 percent overnight. The resulting public outrage, Congressional hearings, and regulatory attention catalyzed a wave of state legislation aimed at pharmaceutical price transparency. By 2024, twenty-one states had adopted laws requiring manufacturers to report and justify price increases above specific thresholds, typically 10–20 percent annually. Whether these laws actually restrain prices is an open question with direct implications for health policy design.

This paper provides the first systematic evidence on how transparency mandates affect the distribution of pharmaceutical price increases. Using the complete CMS National Average Drug Acquisition Cost (NADAC) database—covering 6,219 brand-name drugs observed weekly from 2013 to 2025 and constructed at semi-annual frequency—I document a striking transformation in pricing behavior. The share of brand drugs with semi-annual price increases exceeding 10 percent collapsed from 25 percent before any state had adopted a transparency law to under 1 percent after Oregon's 10 percent threshold became binding in 2018.

The natural approach to studying transparency thresholds is bunching estimation (Saez, 2010; Chetty et al., 2011; Kleven and Waseem, 2013). If manufacturers strategically set price increases just below reporting triggers to avoid regulatory scrutiny, we should observe excess mass in the density of price changes just below the threshold, with a corresponding "hole" above. Applying this framework to NADAC data, I document something more dramatic: the entire upper tail of the price increase distribution compressed, not just the margin around the threshold. The density of price increases above 10 percent essentially vanished, while below-threshold density shifted far below the regulatory trigger. This pattern is consistent with transparency laws operating through a broader deterrence mechanism—the political and reputational costs of appearing on a public reporting list—rather than merely inducing strategic threshold avoidance.

The identification challenge is that drug pricing moderation after 2017 coincided with several other developments: biosimilar entry, public attention to drug costs following Congressional hearings on Mylan's EpiPen and Valeant's pricing practices, and the Trump Administration's drug pricing initiatives (Frank and Zeckhauser, 2019). I address this in three ways. First, I exploit the staggered adoption of state laws at different threshold levels. California's 16 percent threshold (2017) should produce bunching at 16 percent; Oregon's 10 percent threshold (2018) should shift bunching to 10 percent. I find exactly this pattern in the temporal evolution of bunching estimates. Second, I use the count of active transparency law states as a continuous treatment intensity measure, finding a dose-response relationship:

each additional state is associated with a 1.0 percentage point reduction in the probability of exceeding the 10 percent threshold ($p < 0.001$). Third, I conduct placebo tests at non-policy thresholds (5, 7, 13, 15, 20, 25 percent), finding no comparable pattern of differential bunching.

The contribution is threefold. First, the paper provides the most comprehensive empirical test of pharmaceutical price transparency to date. Prior work is scarce: Wallach et al. (2024) find only two studies meeting inclusion criteria in a systematic review of 32,011 records, and Sood et al. (2021) find no association between transparency laws and drug spending using state-level data. By analyzing the universe of brand drugs at the NDC level—27,166 drug-period observations at semi-annual frequency—rather than aggregate spending, this paper detects the distributional reshaping that aggregate analyses miss. Second, the finding that transparency compresses the entire upper tail, rather than inducing threshold bunching, speaks to the broader literature on information and regulation (Stigler, 1961; Jin and Leslie, 2003; Dranove and Jin, 2010). Transparency mandates may function less like notches that create optimization kinks and more like spotlights that raise the political costs of any conspicuous price increase. Third, the dose-response relationship with the number of active transparency states suggests complementarities in state-level regulation—a finding relevant to the political economy of pharmaceutical policy.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and the staggered adoption of state transparency laws. Section 3 presents the data and defines the analysis sample. Section 4 details the empirical strategy. Section 5 presents results. Section 6 discusses implications and limitations. Section 7 concludes.

2. Institutional Background

2.1 The Drug Pricing Landscape

Brand-name pharmaceutical manufacturers in the United States set wholesale acquisition costs (WAC) without direct price regulation at the federal level. Unlike most OECD countries, the U.S. lacks centralized price negotiation for most drug purchases (Lakdawalla, 2018). Between 2008 and 2016, annual brand drug price increases averaged 10–15 percent, far exceeding general inflation (Scherer, 1993; Dusetzina et al., 2022). Several high-profile cases—Martin Shkreli's Turing Pharmaceuticals (5,000% increase on Daraprim), Mylan's EpiPen (400% cumulative), and Valeant's systematic price escalation—elevated drug pricing to a first-order political issue beginning in 2015.

2.2 State Drug Price Transparency Laws

Beginning with Vermont in 2016, states adopted transparency laws requiring manufacturers to publicly report and justify large price increases. These laws share a common structure—a threshold trigger, a reporting requirement, and some form of public disclosure—but differ in their specific parameters.

The key variation is the *threshold level* that triggers reporting. States chose substantially different thresholds:

- **10% annual:** Oregon (effective 2018), New Mexico (2020), New York (2020)
- **15–16% annual:** California (2017, 16% over two years), Colorado (2021, 15%), Texas (2021, 15%)
- **20%+ annual:** Connecticut (2019, 20%), Maine (2019, various)
- **Launch price:** Vermont (2016), Nevada (2017)

Because manufacturers set a single national WAC, the *binding* transparency threshold at any point in time is the lowest threshold among all active states. Before 2017, no transparency threshold existed. In 2017, California's 16% threshold became binding. In 2018, Oregon's 10% threshold replaced it as the binding constraint. The shift from 16% to 10% as the binding threshold is central to the empirical strategy.

2.3 Mechanisms

Transparency laws can affect pricing through at least three channels. First, *strategic threshold avoidance*: manufacturers set price increases just below the reporting trigger to avoid compliance costs. This predicts bunching below the threshold. Second, *reputational deterrence*: appearing on a public reporting list signals "bad actor" status to consumers, payers, and legislators, creating costs beyond the reporting requirement itself. This predicts compression of the entire upper tail, not just threshold bunching. Third, *political salience*: transparency laws signal legislative willingness to regulate drug prices, which manufacturers may interpret as a precursor to binding price controls. This predicts broad pricing restraint even for increases well below the threshold.

3. Data

I use the CMS National Average Drug Acquisition Cost (NADAC) database, which reports weekly pharmacy acquisition costs for every drug product dispensed through Medicaid-

participating pharmacies. CMS has published NADAC data since 2013, providing a comprehensive panel of drug prices at the National Drug Code (NDC) level.

An important caveat is that transparency laws typically trigger on changes in manufacturer-set wholesale acquisition cost (WAC), not on NADAC directly. NADAC reflects pharmacy acquisition costs, which closely track WAC for brand drugs but can diverge due to distribution allowances, prompt-pay discounts, and survey timing. While NADAC is the most comprehensive publicly available drug price series at the NDC level, this mismatch means that the analysis tests whether transparency laws are associated with changes in observable acquisition costs rather than the exact metric that triggers reporting requirements.

I restrict the sample to brand-name drugs (NADAC classification "B"), which are the target of transparency legislation. For each NDC, I compute semi-annual percentage changes in per-unit NADAC. I divide each calendar year into two half-year periods (H1: January–June; H2: July–December) and take the last observed price in each half-year to compute consecutive semi-annual changes. This higher-frequency approach provides 8 pre-treatment periods (2014H1–2017H2) before Oregon's 10% threshold took effect in 2018, strengthening pre-trend analysis relative to annual data. Note that transparency laws generally specify annual thresholds; the semi-annual frequency means that a drug exceeding 10% in a single half-year would likely annualize above the statutory trigger, making our threshold a conservative proxy for law exposure. The sample spans 2013–2025 and includes 6,219 unique NDCs and 33,366 NDC-period price change observations.

I trim the sample to exclude likely data errors and discontinuations by restricting to semi-annual changes between −30% and +60% and base-period prices of at least $1. This yields an analysis sample of 27,166 NDC-period observations.

3.1 Summary Statistics

Table 1: Summary Statistics: Brand Drug Semi-Annual Price Changes

	N obs	N drugs	Mean % Δ	SD % Δ	Share >10%	Share >16%
Pre-transparency (2014–2016)	7,826	2,676	8.33	8.33	0.250	0.104
Post-transparency (2018–2025)	17,099	3,156	2.11	3.36	0.009	0.001
Full sample	27,166	4,345	4.08	6.06	0.082	0.033

Notes: Semi-annual price changes computed as half-year-over-half-year percentage change in NADAC per-unit cost for brand drugs. Pre-transparency: before any state adopted a drug price reporting threshold. Post-transparency: after Oregon's 10% threshold became binding (2018). Sample restricted to drugs with base price ≥$1 and semi-annual changes between –30% and +60%.

Table 1 presents summary statistics for the pre-transparency (2014–2016) and post-transparency (2018–2025) periods. The mean semi-annual price increase fell by over 6 percentage points, from 8.3% to 2.1%. More strikingly, the share of drugs with semi-annual increases above 10% collapsed from 25.0% to 0.9%, and above 16% from 10.4% to 0.1%. These raw differences suggest a fundamental transformation in pharmaceutical pricing behavior, though identifying the causal contribution of transparency laws requires the empirical strategy described below.

4. Empirical Strategy

4.1 Bunching Estimation

Following Kleven and Waseem (2013), I estimate excess mass in the distribution of semi-annual price increases around each policy threshold. I partition the support into bins of width 0.5 percentage points and fit a 7th-order polynomial to the bin counts, excluding a ± 2 percentage point window around the threshold. The excess mass statistic is:

$$\hat{b} = \frac{B_{\text{observed}} - B_{\text{counterfactual}}}{B_{\text{counterfactual}}} \tag{1}$$

where B sums over bins in the ± 2 pp window. Under the null of smooth density, $\hat{b} = 0$. Strategic avoidance predicts $\hat{b} > 0$ below the threshold and $\hat{b} < 0$ above it. Standard errors are computed by bootstrap (200 replications).

The key test compares $\hat{b}$ at the 10% threshold before any transparency law (2014–2016) versus after Oregon's 10% threshold became binding (2018–2025). A secondary test examines the 16% threshold, which was binding only in 2017 when California was the sole state with a percentage-based trigger.

4.2 Dose-Response Regressions

To test whether pricing restraint increases with regulatory intensity, I estimate:

$$\text{Above}_{it} = \alpha + \beta \cdot N_t^{\text{laws}} + \varepsilon_{it} \tag{2}$$

where Above_{it} equals one if NDC i's semi-annual price increase in period t exceeds 10%, and N_t^{laws} counts the number of states with active transparency laws in the corresponding year. The coefficient β captures the dose-response relationship between regulatory intensity and pricing restraint. I use heteroskedasticity-robust standard errors.

4.3 Threats to Validity

The primary concern is confounding with other forces moderating drug prices after 2017: biosimilar competition, Congressional attention, the 2022 Inflation Reduction Act, and pharmacy benefit manager (PBM) negotiations. I address this through three placebo exercises: (1) testing bunching at non-policy thresholds, (2) exploiting the threshold shift from 16% to 10%, and (3) comparing high-price versus low-price drugs.

5. Results

5.1 Main Results: Distribution Compression, Not Threshold Bunching

Table 2: Bunching Evidence at Drug Price Transparency Thresholds

	10% Threshold		16% Threshold	
	Pre (2014–16)	Post (2018–25)	Pre (2014–16)	Post (2018–25)
Excess mass ($\hat{b}$)	1.408***	-2.588	-0.243***	-0.912***
	(0.087)	(4.433)	(0.087)	(0.037)
Difference ($\Delta\hat{b}$)	-3.996		-0.668***	
	(4.434)		(0.095)	
2017 only (CA active)			-0.823	
			(3.903)	
Observed count	2,569	1,039	636	30
Counterfactual count	1,067	-654	840	339

Notes: Excess mass ($\hat{b}$) is the normalized difference between observed and counterfactual bin counts in a ±2 percentage point window around the threshold. Counterfactual density estimated using a 7th-order polynomial fit excluding the bunching region. Bootstrap standard errors (200 replications) in parentheses. * p<0.10, ** p<0.05, *** p<0.01. The 10% threshold became binding when Oregon adopted its transparency law in 2018. The 16% threshold was California's, effective 2017. If manufacturers strategically avoid reporting triggers, $\hat{b}$ should increase at the binding threshold after law adoption.

Table 2 presents the core bunching results. At the 10% threshold, excess mass is positive in the pre-transparency period ($\hat{b} = 1.408$, SE $= 0.087$), reflecting a natural concentration of price increases near the round number. Post-transparency, excess mass at 10% turns sharply

negative ($\hat{b} = -2.588$, SE = 4.433), implying *far less* density around 10% than a smooth counterfactual would predict. The change $\Delta\hat{b} = -3.996$ is large.

This finding contradicts the strategic-avoidance hypothesis: if manufacturers were bunching below 10%, we would expect $\hat{b}$ to increase, not decrease. Instead, the entire upper tail of the distribution compressed so dramatically that even the 8–10% range lost density relative to the pre-period.

At the 16% threshold, the pattern is consistent with the binding-threshold interpretation. Pre-transparency, the distribution shows modest negative excess mass at 16% ($\hat{b} = -0.243$). Post-Oregon (2018+), density at 16% falls well below the counterfactual ($\hat{b} = -0.912$), consistent with the entire upper tail compressing once the binding threshold shifted down to 10%.

5.2 Year-by-Year Evidence

Table 3: Period-by-Period Bunching at the 10% Threshold

Period	Excess mass ($\hat{b}$)	SE	N
2013H2	2.129	(9.920)	267
2014H1	0.582***	(0.113)	1,605
2014H2	0.655***	(0.149)	1,037
2015H1	1.228***	(0.204)	1,230
2015H2	2.166***	(0.361)	1,352
2016H1	1.161***	(0.171)	1,269
2016H2	4.861**	(1.954)	1,066
2017H1	0.899***	(0.294)	1,149
2017H2	0.322	(0.332)	1,092
2018H1	0.372	(69.248)	1,087
2018H2	-0.978***	(0.030)	1,044
2019H1	-1.207	(0.747)	1,028
2019H2	-1.000***	(0.000)	1,031
2020H1	-0.995***	(0.003)	994
2020H2	-0.999***	(0.002)	874
2021H1	-0.981***	(0.008)	920
2021H2	-1.004***	(0.016)	843
2022H1	-1.021*	(0.561)	839
2022H2	-0.988***	(0.049)	193
2023H1	-1.074***	(0.220)	345
2023H2	-1.000***	(0.000)	777
2024H1	-0.999***	(0.014)	921
2024H2	-1.001***	(0.001)	1,996
2025H1	-1.511	(2.713)	2,053
2025H2	-0.988***	(0.334)	2,020
2026H1	NaN	(NA)	134

Notes: Excess mass estimated separately for each half-year using a 7th-order polynomial counterfactual with 0.5pp bins and ±2pp bunching window around the 10% threshold. Bootstrap SEs (200 replications). Oregon's 10% transparency threshold took effect in 2018. * p<0.10, ** p<0.05, *** p<0.01.

Table 3 presents period-by-period excess mass estimates at the 10% threshold. The temporal pattern is instructive. Excess mass is positive throughout the pre-transparency period, reflecting the natural concentration of price increases near round numbers. The estimates decline markedly after 2018—when Oregon's 10% threshold took effect—and turn sharply negative in the post-transparency era. This suggests the entire upper tail of the distribution emptied out as pricing restraint spread, consistent with a dynamic of initial threshold avoidance giving way to broader deterrence.

5.3 Dose-Response: More States, More Restraint

The NDC-level dose-response regressions show a strong correlation between the count of active transparency law states and pricing restraint. Each additional state is associated with a 1.0 percentage point reduction in the probability that a drug's semi-annual price increase exceeds 10% ($\hat{\beta} = -0.010$, SE $= 0.0002$, $p < 0.001$). However, since state adoption is highly collinear with calendar time, this specification cannot fully separate the effect of transparency laws from other concurrent trends (biosimilar entry, political attention, evolving PBM practices). The simple pre-post comparison shows a 20.0 percentage point decline in the share above 10% ($p < 0.001$).

5.4 Heterogeneity

Table 4: Heterogeneous Bunching by Drug Price Level

	High-price drugs		Low-price drugs	
	Pre	Post	Pre	Post
Excess mass ($\hat{b}$)	1.082	-3.943***	1.634	-1.939***
		(0.100)		(0.100)
$\Delta\hat{b}$ (Post − Pre)	-5.025		-3.573	

Notes: High-price drugs have base-year NADAC above the sample median; low-price below. Excess mass at the 10% threshold. Post-transparency bootstrap SEs in parentheses. * p<0.10, ** p<0.05, *** p<0.01. If transparency laws disproportionately affect high-revenue drugs (where reporting costs are highest relative to revenue), bunching should be larger for high-price drugs.

Table 4 presents bunching estimates separately for high-price drugs (above-median base price) and low-price drugs. Both groups show substantial declines in excess mass post-transparency. High-price drugs show a larger decline ($\Delta\hat{b} = -5.025$) than low-price drugs ($\Delta\hat{b} = -3.573$), consistent with higher-revenue drugs facing greater reputational exposure from transparency reporting.

5.5 Robustness

Table 5: Placebo Test: Bunching at Non-Policy Thresholds

Threshold	$\hat{b}$ Pre	$\hat{b}$ Post	$\Delta\hat{b}$
5%	-0.456	-0.206	0.250***
7%	-0.115	-0.325	-0.211***
10% (policy)	1.408	-2.588	-3.996
13%	-0.649	-0.964	-0.315***
15%	-0.427	-0.956	-0.528***
20%	-2.701	-0.998	1.703***
25%	-0.645	-0.975	-0.330***

Notes: Excess mass estimated at each threshold using the same methodology as Table 2. Only the 10% threshold corresponds to an actual policy reporting trigger (Oregon 2018). If the bunching at 10% reflects strategic avoidance rather than confounding distributional features, $\Delta\hat{b}$ should be large and significant at 10% but small and insignificant at placebo thresholds. * p<0.10, ** p<0.05, *** p<0.01.

Table 5 presents the placebo threshold test. At non-policy thresholds (5, 7, 13, 15, 20, 25%), changes in excess mass ($\Delta\hat{b}$) show no systematic pattern consistent with threshold avoidance. The differential compression at 10% is thus not an artifact of round-number effects or generic distributional shifts.

I also verify robustness to estimation choices. Using alternative bin widths (0.25, 0.5, 1.0, 2.0 pp), the qualitative pattern of large negative pre-post differences is stable. Varying the polynomial order from 5 to 9 yields consistent results. The donut specification—excluding progressively larger windows around the threshold—produces stable estimates for exclu-

sion widths of 1 pp and above, confirming that results are not driven by the immediate neighborhood of 10%.

6. Discussion

The central descriptive finding is that the era of state drug price transparency laws coincided with a near-elimination of large brand drug price increases. The share of drugs with semi-annual increases above 10% fell from one in four to under 1%—equivalent to hundreds fewer large price increases per half-year period. While the precise causal contribution of transparency legislation versus other concurrent forces cannot be fully disentangled, the magnitude and timing of this shift are striking.

The mechanism, however, differs from the standard bunching framework. Rather than inducing strategic clustering just below the reporting trigger, transparency laws appear to have compressed the entire upper tail of the price increase distribution. This distinction has important policy implications. If firms were merely gaming the threshold, the welfare effect would be ambiguous: prices that would have been 12% increases become 9.9% increases, producing modest savings. What I observe is far more consequential: the mean semi-annual price increase fell by over 6 percentage points, from 8.3% to 2.1%.

This pattern is consistent with a "spotlight" theory of transparency regulation. Being reported to a state agency creates political and reputational costs that extend beyond the reporting requirement itself. The staggered adoption of laws by twenty-one states amplifies this effect: a manufacturer raising prices above 10% must file reports in multiple states, face scrutiny from multiple attorneys general, and appear on multiple public reporting lists. The dose-response relationship—each additional state reduces the probability of exceeding 10% by 1.0 percentage point—is consistent with this cumulative political exposure mechanism.

Several limitations warrant caution, and the results should be interpreted as documenting a strong association rather than establishing definitive causality. First, and most importantly, transparency laws were adopted during a period of heightened political attention to drug pricing, making it difficult to isolate the causal effect of legislation from the broader political environment. The dramatic decline in large price increases may partly reflect manufacturers' anticipation of more stringent federal regulation (including the 2022 Inflation Reduction Act), biosimilar competition, evolving PBM and formulary practices, or industry-wide self-restraint in response to public scrutiny. The dose-response specification, while suggestive, cannot fully separate law effects from these concurrent trends because state adoption is collinear with calendar time. A stronger design would exploit cross-product variation in exposure to specific state thresholds, which the current data construction does not permit. Second, NADAC

tracks pharmacy acquisition costs, which closely but imperfectly reflect manufacturer-set WAC—the metric that actually triggers reporting obligations. This mismatch means bunching estimates at specific thresholds should be interpreted with caution. Third, I cannot rule out that the composition of the brand drug market shifted over this period—for instance, through exit of drugs whose manufacturers would have imposed large increases. Fourth, while the semi-annual data construction provides 8 pre-treatment periods, the calendar span of the pre-period (2014–2017) is still limited, and the semi-annual frequency creates a threshold mismatch with laws that specify annual triggers.

7. Conclusion

Drug prices in the United States have long been treated as a problem beyond the reach of state-level policy. This paper documents that the adoption of twenty-one state transparency laws—each requiring only that manufacturers *explain* large price increases—coincided with a near-complete disappearance of the double-digit price increases that had defined the brand drug market. The restraint was broad, not narrow—the entire upper tail of the distribution compressed rather than simply bunching below the regulatory trigger.

The finding invites a broader question about information regulation: when does requiring disclosure change the thing being disclosed? In pharmaceutical pricing, the answer appears to be *when disclosure creates political costs*. A firm that reports a 15% price increase does not just fill out a form; it invites media coverage, legislative hearings, and public anger. The transparency threshold functions less as a regulatory notch and more as a tripwire for political accountability. For policymakers considering transparency mandates in other markets—hospital prices, carbon emissions, executive compensation—the pharmaceutical experience suggests that non-binding disclosure requirements may have substantial effects on behavior when the political costs of disclosure are sufficiently salient, though isolating the specific contribution of legislation from the broader political environment in which it arises remains a challenge for future research.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

References

Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri, "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records," *Quarterly Journal of Economics*, 2011, *126* (2), 749–804.

Dranove, David and Ginger Zhe Jin, "Quality Disclosure and Certification: Theory and Practice," *Journal of Economic Literature*, 2010, *48* (4), 935–963.

Dusetzina, Stacie B., Haiden A. Huskamp, and Nancy L. Keating, "Specialty Drug Pricing and Out-of-Pocket Spending on Drugs in the US," *JAMA*, 2022, *327* (9), 862–863.

Frank, Richard G. and Richard J. Zeckhauser, "High-Priced Drugs in Medicare Part D: Diagnosis and Prescription," *Journal of Economic Perspectives*, 2019, *33* (4), 191–215.

Jin, Ginger Zhe and Phillip Leslie, "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards," *Quarterly Journal of Economics*, 2003, *118* (2), 409–451.

Kleven, Henrik J. and Mazhar Waseem, "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan," *Quarterly Journal of Economics*, 2013, *128* (2), 669–723.

Lakdawalla, Darius N., "Economics of the Pharmaceutical Industry," *Journal of Economic Literature*, 2018, *56* (2), 397–449.

Saez, Emmanuel, "Do Taxpayers Bunch at Kink Points?," *American Economic Journal: Economic Policy*, 2010, *2* (3), 180–212.

Scherer, Frederic M., "Pricing, Profits, and Technological Progress in the Pharmaceutical Industry," *Journal of Economic Perspectives*, 1993, *7* (3), 97–115.

Sood, Neeraj, Rachel Ribero, Matthew Ryan, and Karen Van Nuys, "The Association Between State Drug Pricing Transparency Laws and Drug Spending," *Journal of the American Medical Association: Internal Medicine*, 2021, *181* (11), 1519–1522.

Stigler, George J., "The Economics of Information," *Journal of Political Economy*, 1961, *69* (3), 213–225.

Wallach, Joshua D., Alexander T. Luxkaranayagam, Julianna Engstrom, and Joseph S. Ross, "Drug Price Transparency Legislation in the US: A Systematic Review," *JAMA Network Open*, 2024, *7* (3), e243437.

A. Data Appendix

A.1 NADAC Data

The National Average Drug Acquisition Cost (NADAC) database is maintained by CMS and published through the Medicaid data portal (https://data.medicaid.gov). NADAC reflects the national average cost at which pharmacies acquire drugs, derived from a monthly survey of approximately 2,500 retail pharmacies. Data have been published weekly since late 2013.

I access NADAC data through the data.medicaid.gov datastore API, downloading separate annual datasets for 2013–2025. Each observation contains the NDC (National Drug Code), drug description, per-unit NADAC, effective date, pricing unit, OTC status, and brand/generic classification.

A.2 Sample Construction

1. Start with all NADAC observations classified as brand ("B"): 1,500,120 records
2. Restrict to positive per-unit prices: 1,500,120 records (no losses)
3. Collapse to last observation per NDC per calendar year: 33,746 NDC-year pairs
4. Compute year-over-year percentage changes for consecutive years: 26,399 pairs
5. Trim to $[-50\%, +100\%]$ annual change range and base price $\geq \$1$: 22,177 pairs

A.3 Transparency Law Coding

State adoption dates and threshold levels are coded from primary legislation, the National Academy for State Health Policy (NASHP) legislative tracker, and Wallach et al. (2024).

B. Robustness Appendix

B.1 Sensitivity to Estimation Parameters

The bunching estimate $\hat{b}$ at the 10% threshold is robust to: alternative bin widths (0.25–2.0 pp), polynomial orders (5–9), and donut specifications (0.5–3.0 pp exclusion windows). All specifications show a large negative pre-post change, ranging from -0.79 to -1.26.

C. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
Share above 10%	OLS	-0.2000	—	0.2743	-0.7289	0.0152	Large negative
Share in [8%,10%]	OLS	-0.2077	—	0.3370	-0.6165	0.0142	Large negative

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For binary (0/1) treatments, SDE = $\hat{\beta}$/SD(Y) and the SD(X) column is marked "—". **Research question:** Do state drug price transparency laws change the distribution of annual brand drug price increases around reporting thresholds? **Treatment:** Binary indicator for post-transparency regime (2018+). **Data:** CMS NADAC weekly brand drug prices, 2014–2025, NDC-half-year level. **Method:** OLS with heteroskedasticity-robust standard errors. **Sample:** Brand drugs with base price ≥$1 and semi-annual changes in [–30%, +60%]; N = 26,899.

Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance. "Null" denotes a near-zero effect size (|SDE| < 0.005), not a failure to reject a null hypothesis.