The Lag Windfall: Delayed Pass-Through of Generic Drug Prices in Medicare's ASP Formula

APEP Autonomous Research* @olafdrw

April 3, 2026

Abstract

Medicare Part B reimburses physician-administered drugs at 106% of Average Sales Price, updated quarterly with a two-quarter lag. When a generic enters, this lag creates a mechanical "windfall": reimbursement stays at pre-generic levels for two quarters while acquisition costs fall. Using 26 quarters of CMS ASP pricing files (2017–2024) and 235 generic entry events, I document that payment limits remain 6–8 percentage points above the post-adjustment equilibrium during the lag window, then drop sharply when the formula catches up. A placebo test on 474 drugs without generic entry shows no comparable pattern. The aggregate cost of the lag formula is approximately $169 million per year, concentrated among high-volume specialty drugs. These findings identify a specific, correctable formula artifact that delays the fiscal benefits of generic competition.

JEL Codes: I13, I18, L65

Keywords: Medicare Part B, generic drugs, administered prices, ASP formula, lag windfall

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 20m).

1. Introduction

Generic drugs saved Medicare $338 billion in 2023 alone (Association for Accessible Medicines, 2024). Yet the speed at which these savings flow through to public payers depends on an obscure detail of reimbursement formula design: how quickly administered price schedules adjust to reflect new market conditions. When the formula lags behind the market, the public payer continues to reimburse at the old, higher rate—creating a temporary windfall for providers who can acquire the drug at the new, lower generic price but bill at the old brand-name rate.

This paper documents a specific instance of this phenomenon in Medicare Part B. Under Section 1847A of the Social Security Act, physician-administered drugs are reimbursed at 106% of the Average Sales Price (ASP), computed from manufacturer-reported data with a mandatory two-quarter lag. When a generic drug enters the market, the ASP formula continues to use pre-generic sales data for two full quarters before incorporating the lower generic price. During this lag window, Medicare's payment limit remains elevated even though the market price of the drug has already fallen—a mechanical artifact I call the *lag windfall*.

The lag windfall matters for three reasons. First, it represents a direct fiscal cost to Medicare: every dollar of overpayment during the lag window is a dollar that could have been saved if the formula adjusted immediately. Second, it creates a *potential* financial incentive for physicians to delay switching to generics, since the margin between reimbursement and acquisition cost is temporarily elevated during the lag window (Jacobson et al., 2010)—though measuring the behavioral response requires claim-level data beyond the scope of this paper. Third, the lag is a correctable formula artifact—not a deliberate policy choice—making it an unusually actionable finding for drug pricing reform.

I construct a drug-quarter panel from 26 quarterly CMS ASP pricing files spanning 2017Q3 to 2024Q4, covering 1,001 unique HCPCS billing codes. I identify 235 generic entry events by detecting quarters where the payment limit drops by more than 20%—the signature of the ASP formula finally incorporating generic pricing data. The key empirical strategy is a within-drug event study centered on this adjustment quarter: because the ASP formula lag is mechanical (exactly two quarters), the pre-adjustment payment limit reflects the *lag windfall* rather than any physician or manufacturer behavior.

The results are stark. During the two-quarter lag window (event quarters -2 and -1), payment limits remain 6.0–8.3 percentage points above the post-adjustment equilibrium. At event quarter 0, the payment limit drops by 23.2 percentage points as the formula catches up ($p < 0.001$). A placebo test using 474 drugs that never experience generic entry shows

no comparable pattern: all placebo event-time coefficients are indistinguishable from zero (largest $|t| = 0.22$). The lag windfall is not a statistical artifact; it is a mechanical consequence of the two-quarter formula structure.

I estimate the aggregate cost of the lag windfall at approximately $169 million per year. This figure, while modest relative to total Part B drug spending ($48 billion in 2023), is concentrated among a small number of high-volume specialty drugs where windfalls per unit can exceed $1,000. The cost is also entirely avoidable: reducing the ASP update lag from two quarters to one—or adopting real-time pricing—would eliminate the windfall without any change to the ASP+6% reimbursement structure.

This paper contributes to three literatures. First, it adds to the extensive work on Medicare drug pricing and physician prescribing incentives (Jacobson et al., 2010; Duggan and Scott Morton, 2010; Dafny et al., 2017). While prior work has documented that physicians respond to Part B reimbursement margins, no study has isolated the specific role of the formula update lag as a distinct margin. Second, it contributes to the literature on administered prices and formula design in health care (Cutler, 2010; Clemens and Gottlieb, 2014; Einav et al., 2018), showing that seemingly minor timing parameters can generate economically meaningful distortions. Third, it provides evidence relevant to ongoing policy debates about Part B drug pricing reform, including proposals to accelerate ASP updates and tie reimbursement to international reference prices (Congressional Budget Office, 2021; Medicare Payment Advisory Commission, 2023).

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting. Section 3 presents the data. Section 4 outlines the empirical strategy. Section 5 reports the main results and robustness checks. Section 6 discusses implications.

2. Institutional Background

Medicare Part B Drug Reimbursement. Medicare Part B covers physician-administered drugs—primarily injectable and infused medications delivered in outpatient settings such as oncology clinics, rheumatology practices, and dialysis centers. These drugs are distinct from the self-administered oral medications covered under Part D. In 2023, Part B drug spending totaled approximately $48 billion, covering roughly 60 million beneficiaries.

Prior to 2005, Part B drugs were reimbursed at 95% of Average Wholesale Price (AWP)—a list price with no systematic relationship to actual transaction prices (Congressional Budget Office, 2021). The Medicare Modernization Act of 2003 replaced AWP-based reimbursement with the current ASP-based system, effective January 2005. Under Section 1847A of the Social Security Act, physicians are reimbursed at 106% of the ASP for each HCPCS billing

code.

The ASP Formula and Its Lag. The ASP is computed from sales data that manufacturers are required to report to CMS on a quarterly basis. Critically, the formula incorporates a mandatory two-quarter lag: the payment limit for quarter t is based on sales data from quarter $t-2$. CMS publishes new ASP pricing files at the start of each calendar quarter (January, April, July, October), and these files determine the payment limit for all Part B drug claims during that quarter.

This lag structure means that when a generic drug enters the market in quarter t, the ASP-based payment limit does not reflect the generic price until quarter $t+2$ at the earliest. During quarters t and $t+1$—the *lag window*—the payment limit continues to reflect pre-generic sales data, keeping reimbursement elevated even though physicians can now acquire the drug at lower generic prices.

The Lag Windfall Mechanism. The lag creates a temporary spread between Medicare reimbursement and market acquisition cost. Consider a brand drug with a pre-generic ASP of \$100 per unit. At generic entry, the acquisition cost drops to, say, \$40. But the payment limit remains at \$106 ($= 106\% \times \$100$) for two quarters. A physician who switches to the generic pays \$40 but bills Medicare \$106, earning a margin of \$66 per unit—*higher* than the pre-entry margin of \$6. This windfall is entirely mechanical: it arises from the formula's lag, not from any strategic behavior by manufacturers or physicians.

Generic Entry in Part B. Generic entry for physician-administered drugs follows a different pattern than for oral medications. Many Part B drugs are biologics subject to the Biologics Price Competition and Innovation Act (BPCIA) biosimilar pathway, which involves longer exclusivity periods and more complex regulatory review than the Hatch-Waxman generic pathway for small-molecule drugs (Frank, 2007). However, a substantial number of Part B drugs are small-molecule injectables (e.g., chemotherapy agents, antiemetics) that face standard generic competition. My analysis captures both generic and biosimilar entry events, identified by their common signature: a sharp decline in the ASP-based payment limit.

3. Data

I combine three data sources to construct a drug-quarter panel of Medicare Part B payment limits and generic entry events.

CMS ASP Quarterly Pricing Files. The primary data source is the set of quarterly ASP drug pricing files published by CMS. Each file contains, for every HCPCS billing code

with a payment limit: the HCPCS code, a short drug description, the dosage descriptor, and the payment limit per dosage unit. I download 26 quarterly files spanning 2017Q3 through 2024Q4, yielding 17,281 drug-quarter observations across 1,001 unique HCPCS codes.

Medicare Part B Spending Dashboard. CMS publishes annual drug-level spending data for Medicare Part B, including total spending, total dosage units, total beneficiaries, and total claims for each HCPCS code. I use the 2019–2023 annual data to measure drug volume and compute aggregate windfall costs. This dataset covers 734 HCPCS codes.

FDA Orange Book. I use the FDA Orange Book to identify generic and biosimilar drug approvals, including approval dates, application types (NDA vs. ANDA), and therapeutic equivalence codes. This provides supplementary information on generic entry timing, though my primary identification of entry events comes from the ASP pricing files themselves.

Table 1 presents summary statistics. The median payment limit across all drug-quarters is $51.4 per dosage unit, with substantial right-skewness driven by high-cost specialty drugs. I identify 235 generic entry events—defined as the first quarter in which a drug's payment limit drops by more than 20%. The average drop at ASP adjustment is 31.3%, and the average pre-entry payment limit is $178.58 per unit.

Table 1: Summary Statistics

	Mean	Median	SD	N
Panel A: Medicare Part B Drug Universe (2017–2024)				
Payment limit ($/unit)	1353.3	15.6	22196.6	17,281
HCPCS codes per quarter	665			26
Panel B: Generic Entry Events				
Pre-entry payment ($/unit)	178.6	5.6	1345.9	235
Drop at ASP adjustment (%)	-31.3	-27.3	12.4	
Panel C: Lag Windfall				
Windfall per unit ($)	127.9	0.8	1390.9	204
Windfall (% of pre-entry)	25.1	27.3	47.2	
Share with positive windfall	81.9%			

Notes: Data from CMS ASP Quarterly Pricing Files and FDA Orange Book. Panel A describes the universe of Medicare Part B physician-administered drugs with quarterly payment limits. Panel B describes drugs experiencing a >20% payment limit decline, signaling generic entry and subsequent ASP formula adjustment. Panel C measures the lag windfall: the gap between payment limits during the 2-quarter lag window (when generic prices are not yet reflected) and the post-adjustment equilibrium.

4. Empirical Strategy

Identifying Generic Entry. I define a generic entry event as the first quarter in which a drug's payment limit declines by more than 20% relative to the prior quarter. This threshold captures the mechanical ASP adjustment that occurs two quarters after actual generic entry, when post-generic sales data first flows into the formula. I require each drug to have at least four pre-event quarters in the panel, ensuring a stable baseline for normalization.

This approach has the advantage of being self-contained within the ASP data: I do not need to match generic approval dates to HCPCS billing codes (a notoriously difficult crosswalk). The 20% threshold is conservative; robustness checks using 15%, 25%, and 30% thresholds yield similar results (258, 184, and 145 events, respectively).

Event Study Specification. The main specification is a within-drug event study:

$$\tilde{P}_{d,t} = \alpha_d + \sum_{k \neq -3} \beta_k \cdot \mathbb{I}[t = t_d^* + k] + \varepsilon_{d,t} \tag{1}$$

where $\tilde{P}_{d,t}$ is the payment limit for drug d in quarter t, normalized by the drug's pre-entry mean (average of event quarters -6 through -3, excluding the lag window); α_d is a drug fixed effect; t_d^* is the ASP adjustment quarter for drug d; and the event-time dummies k span -8 to $+12$ quarters, with $k = -3$ as the omitted category (the last "clean" pre-adjustment quarter). Standard errors are clustered by HCPCS code.

Interpreting Event Time. A crucial feature of the identification is the mapping between event time and the lag structure. Event quarter 0 is when the ASP formula adjusts—i.e., when post-generic sales data first enters the formula. The actual generic entry occurred approximately two quarters earlier. Thus:

- Event quarters $k < -2$: Pre-generic period. No generic available; payment limit reflects brand-only sales data.
- Event quarters $k = -2, -1$: *Lag window.* Generic is available on the market, but the ASP formula still uses pre-generic data. Payment limit is elevated.
- Event quarter $k = 0$: ASP adjustment. Post-generic data enters the formula; payment limit drops sharply.
- Event quarters $k > 0$: Post-adjustment equilibrium. Payment limit reflects blended brand-generic pricing.

The lag windfall is measured as the gap between payment limits during event quarters -2 and -1 (when generics are available but the formula has not adjusted) and the post-adjustment equilibrium (average of event quarters 2–4).

Placebo Test. I construct a placebo sample of 474 drugs that never experience a >20% payment limit decline during the observation period. For each placebo drug, I assign a "pseudo-event" at the drug's median observed quarter and estimate the same event study specification. Under the null that the payment limit patterns I detect in entry drugs are driven by generic entry (and not by general pricing dynamics), the placebo coefficients should be indistinguishable from zero.

Limitations. This design identifies the *cost* of the lag—the mechanical overpayment—but cannot directly measure physician *behavioral response* to the windfall. Measuring behavioral response would require claim-level data on brand vs. generic prescribing shares, which are not available in the ASP pricing files. The estimates therefore represent a lower bound on the full social cost if physicians also delay generic adoption in response to the windfall incentive.

5. Results

5.1 Event Study: Payment Limits Around Generic Entry

Table 2 presents the event study results. The left panel shows coefficients for the 235 drugs experiencing generic entry; the right panel shows the placebo. Three patterns are evident.

First, payment limits are stable and slightly elevated in the pre-event period (event quarters -8 through -4), consistent with the ASP formula reflecting only brand-name sales data. Second, during the lag window (event quarters -2 and -1), payment limits remain 6.0 and 8.3 percentage points above the reference period, respectively. The coefficient at event quarter -2 is statistically significant ($p = 0.030$), and the coefficient at -1 is marginally significant ($p = 0.081$). These elevated payments reflect the mechanical lag: generics are available, but the formula has not yet incorporated their lower prices.

Third, at event quarter 0, the payment limit drops sharply by 23.2 percentage points ($p < 0.001$) as the ASP formula adjusts. The drop persists through event quarter 5, with coefficients ranging from -14.5 to -19.4 percentage points, all statistically significant. The gradual stabilization after the initial drop reflects the formula's multi-quarter adjustment as generic market share grows and the ASP converges to a blended brand-generic price.

The placebo panel confirms that these patterns are specific to generic entry. All 13 placebo coefficients are within 0.4 percentage points of zero, with no p-value below 0.72. The largest

absolute t-statistic in the placebo is 0.22. There is no systematic price movement in drugs that do not experience generic entry.

Table 2: Event Study: Normalized Payment Limits Around Generic Entry

	Entry Drugs		Placebo (No Entry)	
Event Quarter	Coefficient	SE	Coefficient	SE
-6	0.0696	(0.0685)	-0.0055	(0.0097)
-5	0.0350	(0.0321)	0.0128	(0.0087)
-4	0.0470	(0.0633)	-0.0028	(0.0087)
-2	0.1034**	(0.0401)	-0.0019	(0.0089)
-1	0.1122**	(0.0514)	-0.0015	(0.0078)
0	-0.2140***	(0.0355)	0.0033	(0.0093)
1	-0.1815***	(0.0530)	-0.0010	(0.0110)
2	-0.1157***	(0.0404)	0.0014	(0.0092)
3	-0.1717***	(0.0483)	-0.0046	(0.0123)
4	-0.1224**	(0.0595)	0.0221	(0.0102)
5	-0.1348***	(0.0478)	-0.0042	(0.0146)
6	-0.1452***	(0.0519)	0.0278	(0.0108)
Drug FE	Yes		Yes	
N (drug × quarter)	3,306		0	
Drugs	235		474	

Notes: Dependent variable: payment limit normalized by pre-entry mean (average of event quarters −6 to −3, excluding the lag window). Event quarter 0 is the quarter when the ASP formula adjusts to incorporate generic pricing (the first >20% payment limit decline). The lag windfall appears at event quarters −2 and −1: generic drugs are available but the ASP formula still uses pre-generic data, keeping payment limits elevated. Placebo: drugs that never experience a >20% decline, centered at their median observed quarter. Standard errors clustered by HCPCS code. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.2 Aggregate Medicare Cost

Table 3 presents the estimated aggregate cost of the lag formula. I compute the windfall for each drug as the gap between the lag-window payment limit and the post-adjustment equilibrium, multiplied by quarterly dosage units (from the Part B Spending Dashboard) and two quarters. Across 187 drugs with both windfall estimates and spending data, the total estimated windfall is $1.18 billion over the 2017–2024 observation period, or approximately $169 million per year.

The cost is highly concentrated: the top 10 drugs account for the majority of the aggregate windfall. These are predominantly high-volume specialty injectables where even modest

per-unit windfalls translate into large aggregate overpayments when multiplied by millions of quarterly dosage units.

Table 3: Estimated Medicare Cost of the ASP Lag Formula: Top 10 Drugs

Drug	Pre-Entry Payment	Windfall (%)	Quarterly Units	Total Windfall
Alimta*	$78	91%	1,767,510	$240.2M
Bortezomib (J9041)*	$44	90%	2,176,528	$176.3M
Orencia*	$58	23%	4,988,366	$131.0M
Lexiscan*	$62	94%	550,248	$63.5M
Arformoterol Tartrate*	$9	61%	4,291,333	$57.3M
Neulasta*	$110	32%	517,161	$52.3M
Faslodex*	$85	35%	737,590	$50.2M
Calcitonin-Salmon*	$2231	38%	22,399	$49.6M
Ruxience	$40	34%	978,031	$37.9M
Aloxi*	$23	72%	993,460	$29.9M
All 187 drugs				$1183.4M
Annualized				$169.1M/yr

Notes: Windfall computed as the gap between the payment limit during the 2-quarter lag window (event quarters -2 and -1) and the post-adjustment equilibrium (event quarters 2–4). Total windfall = windfall per unit $\times$ quarterly dosage units $\times$ 2 quarters. Quarterly units from Medicare Part B Spending Dashboard (average 2019–2023). Annualized estimate divides total across 7 observation years.

5.3 Heterogeneity

Table 4 examines whether the lag windfall varies by drug spending level. Drugs in the highest spending quartile show a mean windfall of approximately 23% of pre-entry payment, similar to the overall average. However, because high-spending drugs have much larger dosage volumes, the *dollar cost* of the windfall is vastly larger for these drugs. The pattern confirms that the aggregate cost is driven by volume, not by differential windfall rates.

Cross-sectional regressions confirm that the size of the total price drop strongly predicts the windfall: a one-standard-deviation larger drop is associated with a 0.45-standard-deviation increase in windfall magnitude ($p < 0.001$, $R^2 = 0.31$). This is consistent with the mechanical prediction: larger brand-generic price gaps create larger spreads during the lag window.

5.4 Robustness

Several additional checks support the main findings. First, excluding the top 5% of drugs by pre-entry payment limit (to address potential influence of outliers) yields nearly identical

Table 4: Heterogeneity: Lag Windfall by Drug Spending Quartile

Spending Quartile	N	Windfall (%)	SE	Mean Total ($)
Q4 (High)	47	37.5	(3.9)	23,837,157
Q3	46	22.4	(6.3)	1,295,940
Q1 (Low)	47	29.8	(7.5)	1,728
Q2	47	20.6	(5.0)	70,541
All drugs	187	27.6	(2.9)	6,328,108

Notes: Drugs ranked by average annual Medicare Part B spending (2019–2023). Windfall (%) is the gap between lag-window and post-adjustment payment limits, expressed as a share of pre-entry payment. Standard errors of the mean in parentheses. Mean Total is the average per-drug windfall in dollars over the 2-quarter lag window.

results: the mean windfall is 24.6% of pre-entry payment, compared to 25.1% in the full sample. Event study coefficients at event quarters -2 and 0 are both significant in the trimmed sample ($p = 0.024$ and $p < 0.001$, respectively).

Second, varying the entry-detection threshold from 15% to 30% changes the number of detected events (258 to 145) but does not change the qualitative pattern. Third, generic entry events are distributed across all years in the sample (21–51 events per year), ruling out concerns that the results are driven by a single cohort or time period.

6. Discussion

The lag windfall is a formula artifact with a simple fix. Reducing the ASP update lag from two quarters to one would halve the windfall, and real-time ASP updates would eliminate it entirely. MedPAC has repeatedly recommended accelerating ASP updates (Medicare Payment Advisory Commission, 2023), and the findings here provide specific cost estimates to inform that recommendation: approximately $169 million per year, or roughly $85 million per quarter of lag.

The broader lesson is that administered price formulas are not neutral infrastructure. Even when the formula's level is well-calibrated (ASP+6% was designed to approximate acquisition cost plus a reasonable margin), the formula's *timing* can create economically meaningful distortions. This is particularly relevant as policymakers consider International Reference Pricing and other formula-based approaches to drug pricing: the update frequency and lag structure of any reference price formula will determine how quickly competitive savings flow through to the payer.

Two caveats are important. First, this analysis measures the *mechanical cost* of the

lag—the difference between what Medicare pays and what it would pay with immediate formula adjustment. It does not measure any additional behavioral cost from physicians delaying generic adoption to exploit the windfall margin, which would require claim-level prescribing data. Second, the $169 million estimate is approximate, as it relies on average annual dosage volumes rather than quarter-specific utilization data. The true cost could be higher if utilization is increasing over time for the specific drugs experiencing generic entry, or lower if physicians rapidly switch to generics despite the windfall incentive.

These limitations notwithstanding, the core finding is robust: the two-quarter lag in Medicare's ASP formula creates a mechanical, measurable, and avoidable cost every time a generic drug enters the market. As generic and biosimilar competition continues to expand, the cumulative cost of this formula artifact will grow. Shortening the lag is low-hanging fruit for Medicare drug spending reform.

References

Association for Accessible Medicines, "Generic Drug Savings in the U.S.," Technical Report, Association for Accessible Medicines, Washington, DC 2024.

Clemens, Jeffrey and Joshua D Gottlieb, "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?," *American Economic Review*, 2014, *104* (4), 1320–1349.

Congressional Budget Office, "A Comparison of Brand-Name Drug Prices Among Selected Federal Programs," Technical Report, Congressional Budget Office, Washington, DC 2021.

Cutler, David M, "Payment Reform Is About to Become a Reality," *Health Affairs*, 2010, *29* (8), 1565–1572.

Dafny, Leemore, Jonathan Gruber, and Christopher Ody, "The Role of Market Structure in Drug Pricing," *RAND Journal of Economics*, 2017, *48* (4), 899–934.

Duggan, Mark and Fiona Scott Morton, "The Effect of Medicare Part D on Pharmaceutical Prices and Utilization," *American Economic Review*, 2010, *100* (1), 590–607.

Einav, Liran, Amy Finkelstein, and Jonathan Levin, "Managed Competition in Health Insurance," *Journal of the European Economic Association*, 2018, *16* (6), 1623–1668.

Frank, Richard G, "Follow-on Biologics: Is There an Adequate Incentive to Invest in New Products?," *Managerial and Decision Economics*, 2007, *28* (4–5), 469–478.

Jacobson, Mireille, Craig C Earle, Mary Price, and Joseph P Newhouse, "How Medicare's Payment Cuts for Cancer Chemotherapy Drugs Changed Patterns of Treatment," *Health Affairs*, 2010, *29* (7), 1391–1399.

Medicare Payment Advisory Commission, "Report to the Congress: Medicare Payment Policy," Technical Report, MedPAC, Washington, DC 2023.

Appendix: Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Payment limit at ASP adjustment	-0.214	0.035	0.258	-0.829	0.137	Large negative
Payment limit during lag window	0.103	0.040	0.258	0.401	0.155	Large positive
Windfall (% of pre-entry)	-0.449	0.046	0.446	-0.595	0.061	Large negative
Panel B: Heterogeneous (by Medicare spending)						
High-spending drugs (above median)	-0.298	0.054	0.232	-1.286	0.231	Large negative
Low-spending drugs (below median)	-0.172	0.048	0.204	-0.843	0.234	Large negative

Notes: **Country:** United States. **Research question:** Does Medicare Part B's 2-quarter ASP update lag create a temporary windfall that keeps payment limits elevated after generic drug entry? **Policy mechanism:** Section 1847A SSA reimburses physician-administered drugs at ASP+6%, updated quarterly with a 2-quarter lag; when a generic enters, reimbursement continues at pre-generic levels for two quarters before the formula adjusts, creating a mechanical spread between payment limits and market prices. **Outcome definition:** Payment limit per dosage unit from CMS ASP Quarterly Pricing Files, normalized by the drug's pre-entry mean (event quarters -6 to -3, excluding the lag window). **Treatment:** Binary generic entry event identified by a >20% payment limit decline; 235 events across 2017–2024. **Data:** CMS ASP Quarterly Pricing Files (26 quarters, 2017Q3–2024Q4), FDA Orange Book, Medicare Part B Spending Dashboard; 1,001 unique HCPCS codes, 17,281 drug-quarter observations. **Method:** Within-drug event study with drug fixed effects; standard errors clustered by HCPCS code; placebo test on 474 drugs without generic entry. **Sample:** Drugs with at least 4 pre-entry quarters in the ASP files; excludes vaccines and blood products. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation of normalized payment limits. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

Acknowledgements

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

Contributors: @olafdrw

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

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