The Disclosure Cliff: How Reporting Thresholds Censor Pharmaceutical Payment Data

APEP Autonomous Research* @ai1scl

March 30, 2026

Abstract

The Physician Payments Sunshine Act exempts individual payments below a CPI-indexed per-transaction threshold from mandatory disclosure unless they cumulate above a separate annual aggregate. Using bunching estimation on 433,131 CMS Open Payments records (2018–2024), we quantify the informational blind spot this creates. The payment distribution exhibits significant missing mass just below the threshold (pooled $\hat{b} = -1.11$, SE $= 0.25$), indicating that the density of observed payments drops sharply where the per-transaction reporting rule loses force. The censoring gap is concentrated entirely in food and beverage payments, consistent with elastic payment sizing; non-food categories are too sparse in this range to estimate precisely. This blind spot means the public database systematically understates the frequency of small manufacturer-physician financial interactions—precisely the informal contacts most closely linked to prescribing influence.

JEL Codes: I11, I18, L65, K32

Keywords: pharmaceutical payments, Sunshine Act, bunching estimation, reporting threshold, transparency regulation, censoring

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 2h 5m).

1. Introduction

Every year, pharmaceutical manufacturers buy millions of meals for the physicians who prescribe their products. A substantial literature documents that these small, frequent transfers of value predict prescribing behavior: physicians who receive industry lunches prescribe the sponsor's brand-name drugs at higher rates (**?**), adopt new products faster (**?**), and reduce generic substitution (**?**). The public database that makes this research possible—CMS Open Payments—is widely treated as a near-census of manufacturer-physician financial relationships.

It is not. The Physician Payments Sunshine Act, which mandates the database, exempts individual payments below a per-transaction floor from mandatory reporting. Unless the annual cumulative total to a single physician exceeds a separate aggregate threshold, sub-floor payments simply do not appear in the data. Between 2018 and 2024, this floor has risen from \$10.42 to \$12.70, tracking the consumer price index. The question this paper asks is: how large is the informational blind spot this threshold creates?

We answer this question using a bunching estimator (**??**). The estimator fits a smooth polynomial to the distribution of payment amounts away from the threshold, extrapolates the counterfactual density into the threshold region, and measures any excess or missing mass. If the threshold simply censors some payments from the data—because they fall below the per-transaction minimum and their annual aggregate is too small to trigger the separate reporting rule—the observed distribution should exhibit a density discontinuity at the threshold: fewer payments just below than the smooth counterfactual predicts.

That is precisely what we find. Across four program years (2018, 2020, 2022, 2024) and 433,131 payment records in the \$2–\$30 range, the distribution exhibits significant missing mass just below the reporting threshold. The pooled bunching estimate is $\hat{b} = -1.11$ (SE = 0.25), meaning the observed density in the sub-threshold manipulation region is substantially below the polynomial counterfactual. Year-by-year estimates are consistently negative, ranging from -0.65 to -2.63. The censoring gap tracks the CPI-adjusted threshold as it shifts across years, confirming that it is driven by the reporting rule rather than by round-number pricing or other mechanical features of the payment distribution.

The heterogeneity results sharpen the interpretation. Food and beverage payments—which account for 97 percent of records in our analytic window and whose amounts are trivially adjustable—drive the entire censoring pattern. Non-food payment types are too sparse in the \$2–\$30 range to produce reliable estimates. This is consistent with the institutional mechanism: a pharmaceutical sales representative can easily size a physician lunch to fall below the reporting floor, causing it to vanish from the public record unless the cumulative annual relationship triggers the aggregate rule. Consulting fees, speaker honoraria, and other contractually specified payments are less elastic and less susceptible to threshold-driven censoring.

This paper makes three contributions. First, we provide the first quantitative estimate of the informational blind spot in CMS Open Payments data. Researchers using this database to study industry-physician relationships (**????**) should be aware that small, frequent payments—arguably the most consequential for relationship maintenance—are systematically underrepresented. Second, we demonstrate that the censoring gap is policy-driven: it tracks the CPI-adjusted threshold across years, which would not be the case if the missing mass reflected only the natural thinning of the payment distribution at lower values. Third, we contribute to the disclosure regulation literature (**???**) by documenting a mechanism through which threshold-based reporting rules create blind spots that limit the informational value of mandated databases.

The rest of the paper proceeds as follows. Section 2 describes the institutional background. Section 3 presents the empirical strategy. Section 4 describes the data. Section 5 reports results. Section 6 discusses robustness. Section 7 concludes.

2. Institutional Background

The Physician Payments Sunshine Act. Section 6002 of the Affordable Care Act requires manufacturers of drugs, devices, and biologics to report to CMS all payments or transfers of value made to covered physicians. CMS publishes these records annually in the Open Payments database (**?**). The law was designed to illuminate financial relationships that prior research linked to prescribing behavior (**??**), and has generated a dataset containing over 100 million records spanning more than a decade.

The per-transaction threshold. The Act exempts individual payments below a CPI-indexed per-transaction minimum from mandatory reporting. CMS adjusts this floor annually using the CPI-U and announces it in the Federal Register. Table 1 shows the threshold for each program year from 2018 to 2024. The threshold rose from \$10.42 to \$12.70—a 22 percent increase over seven years.

Table 1: CPI-Indexed Per-Transaction Reporting Thresholds

Program Year	Threshold ($)
2018	10.42
2020	10.73
2022	11.29
2024	12.70

Notes: From annual CMS Federal Register notices. Payments below the threshold are excluded from Open Payments unless annual aggregate payments to the same physician exceed approximately $100 (also CPI-adjusted).

The aggregate reporting rule. A second rule requires reporting of sub-threshold payments if the manufacturer's cumulative annual payments to a single physician exceed a separate aggregate ceiling (approximately $100, also CPI-adjusted). This means that for manufacturer-physician pairs with intensive relationships, even below-threshold payments are reported. The per-transaction threshold primarily censors payments from *low-frequency* relationships—one-off or occasional interactions where the annual total stays below $100.

Food and beverage payments. Food and beverage payments dominate the small-payment distribution, accounting for the overwhelming majority of records between $5 and $30. A typical transaction is a pharmaceutical sales representative providing lunch for a physician's office staff during a product presentation. These amounts are highly elastic: the representative chooses the venue, menu, and per-person cost. This makes food and beverage the category most susceptible to threshold-driven censoring—a representative whose standard lunch costs $12 per physician can, in a year when the threshold is $11.84, simply choose a less expensive option and cause the payment to vanish from the database.

Enforcement. Violations carry civil monetary penalties up to $100,000 for knowing failures to report (**?**). However, keeping a payment below the threshold is not a violation—it is compliance. The enforcement landscape implies that threshold-driven censoring carries zero legal risk.

3. Empirical Strategy

The bunching estimator. We use the bunching framework of **?** and **?**. For each program year t with threshold τ_t, we construct a histogram of payment amounts in \$0.25 bins over a bandwidth $[\tau_t - 3, \tau_t + 3]$. We define an exclusion region $[\tau_t - 1.5, \tau_t + 0.5]$ around the threshold, fit a 7th-degree polynomial to the bin counts outside this region, and extrapolate the counterfactual density into the exclusion zone.

The bunching mass is:

$$\hat{B} = \sum_{j \in \text{below}} \left(c_j - \hat{c}_j^0 \right) \tag{1}$$

where c_j is the observed count in bin j and $\hat{c}_j^0$ is the counterfactual. We normalize by the counterfactual density at the threshold:

$$\hat{b} = \frac{\hat{B}}{\hat{c}^0(\tau_t)} \tag{2}$$

Negative $\hat{b}$ indicates missing mass—fewer payments below the threshold than the counterfactual predicts. Standard errors are computed via 200 Poisson bootstrap replications.

Pooled estimation. For the pooled estimate, we center each year's distribution on its own threshold (computing distance from threshold for each payment), aggregate the centered bins across years, and apply the bunching estimator to the pooled distribution.

Cross-year identification. The CPI adjustment provides a natural validation. If the density discontinuity is driven by the reporting threshold, it should shift with the threshold across years. If it instead reflects round-number pricing or other fixed features of the payment distribution, it should remain anchored at a fixed dollar amount regardless of the threshold.

Heterogeneity. We split the sample into food and beverage versus all other payment types and estimate the bunching separately. Under the censoring mechanism, the gap should be larger for food and beverage payments, where amounts are more elastic and more easily sized below the threshold.

4. Data

Source. We use CMS Open Payments General Payment Data for program years 2018, 2020, 2022, and 2024, accessed via the DKAN datastore API at `openpaymentsdata.cms.gov` (**?**). For each year, we retrieved up to 160,000 payment records via paginated API queries (320

Table 2: Summary Statistics: CMS Open Payments, \$2–\$30 Range

Year	N	Mean	SD	Median	P10	P90	% Food	Threshold
2018	113,633	14.57	5.44	14.29	7.70	21.73	97.4	10.42
2020	113,521	15.67	5.26	15.25	9.83	22.75	95.3	10.73
2022	120,181	16.44	5.71	16.51	8.16	23.73	99.3	11.29
2024	85,796	17.88	6.64	18.42	7.41	26.30	93.1	12.70

Notes: Each row summarizes CMS Open Payments general payment records with amounts between \$2 and \$30 for the indicated program year. % Food is the share classified as Food and Beverage. Threshold is the CPI-indexed per-transaction reporting minimum published in the Federal Register.

pages of 500 records each) and filtered to the \$2–\$30 range, yielding 433,131 records. The API returns records in an order determined by the database (not sorted by payment amount), and our sample spans diverse manufacturers, geographies, and payment dates within each year. Nonetheless, the sample represents approximately 1 percent of each year's full dataset, and we discuss the implications of this limitation below.

Sample composition. Table 2 reports summary statistics. The median payment rises from \$14.29 in 2018 to \$18.42 in 2024, tracking the rightward shift in the reporting threshold. Food and beverage payments account for 96.6 percent of records in the analytic window. The remaining 3.4 percent comprise education (1.2%), travel (1.1%), consulting fees (0.8%), and other categories. Round-dollar heaping is modest: only 3.9 percent of payments fall on exact dollar amounts.

5. Results

Year-by-year estimates. Table 3 reports the bunching estimates. In every program year, $\hat{b}$ is negative and economically large: -1.46 (2018), -2.63 (2020), -0.65 (2022), and -2.58 (2024). The estimates indicate that the observed density of payments in the sub-threshold manipulation region is substantially below the smooth counterfactual. The associated excess mass counts—the estimated number of "missing" payments in the manipulation region—range from roughly 700 to 4,900 per year.

Interpreting the pooled estimate. The pooled estimate of $\hat{b} = -1.11$ (SE = 0.25) confirms that the censoring discontinuity is a robust feature of the data, not an artifact of any single year. To put this in context: the counterfactual polynomial predicts a certain density of payments in the region just below the threshold, and the observed data has roughly half that

Table 3: Bunching at the Reporting Threshold by Program Year

Year	Threshold	Excess Mass	$\hat{b}$	SE($\hat{b}$)	Missing Mass
2018	$10.42	-2,936	-1.455	0.237	-210
2020	$10.73	-4,900	-2.631	0.215	1,180
2022	$11.29	-740	-0.645	0.450	116
2024	$12.70	-2,890	-2.580	0.254	1,159
Pooled	—	-5,775	-1.108	0.249	-805

Notes: Bunching estimates from a polynomial counterfactual (order 7) fitted to $0.25 bins in $[\tau - 3,\ \tau + 3]$, excluding the manipulation region $[\tau - 1.5,\ \tau + 0.5]$. $\hat{b}$ is the normalized excess mass (excess bunching / counterfactual density at threshold). Standard errors from 200 Poisson bootstrap replications. Pooled estimate centers each year's distribution on its own threshold before aggregating.

Table 4: Bunching Heterogeneity: Food & Beverage vs. Other Payment Types

Payment Type	$\hat{b}$	SE($\hat{b}$)	Share of Records (%)
Food & Beverage	-1.580	0.202	96.6
All Other Types	10.757	262.498	3.4

Notes: Bunching estimates by payment type, pooled across years. Food & Beverage payments have highly elastic amounts (meals can be sized to any value), while other types (consulting fees, education, compensation) tend to be set at negotiated round numbers. Same polynomial specification as Table 3.

density. The missing payments are those from manufacturer-physician pairs whose annual aggregate stayed below the $100 aggregate threshold—occasional, low-intensity interactions that the per-transaction exemption allows to escape the database entirely.

Heterogeneity by payment type. Table 4 presents the food-and-beverage versus other-payments split. Food and beverage payments exhibit $\hat{b} = -1.58$ (SE = 0.20), while other payment types are too noisy to estimate precisely. This concentration in the elastic category is exactly what the censoring mechanism predicts: food and beverage amounts are trivially adjustable, making it easy for representatives to keep per-person costs below the threshold. Consulting and speaking fees are contractually set and rarely fall in the $5–$25 range.

Table 5: Robustness of Bunching Estimates

Specification	Variation	$\hat{b}$	SE($\hat{b}$)
Panel A: Polynomial Order			
	Order 5	-1.977	0.081
	Order 6	-2.107	0.082
	Order 7 (baseline)	-1.108	0.268
	Order 8	-1.215	0.257
	Order 9	-1.333	0.592
Panel B: Bandwidth			
	$2	-55.153	208556.880
	$3 (baseline)	-1.108	0.243
	$4	-2.100	0.066
	$5	-2.235	0.047
Panel C: Exclusion Region (below threshold)			
	$0.5	-0.742	0.029
	$1.0	-0.990	0.074
	$1.5 (baseline)	-1.108	0.243
	$2.0	0.196	0.914

Notes: Sensitivity of the pooled bunching estimate to polynomial order (Panel A), analysis bandwidth (Panel B), and the width of the excluded manipulation region below the threshold (Panel C). Baseline specification: 7th-order polynomial, $3 bandwidth, exclusion $[\tau - 1.5, \tau + 0.5]$.

6. Robustness

Specification sensitivity. Table 5 reports robustness across polynomial orders (5–9), bandwidths ($2–$5), and exclusion region widths ($0.50–$2.00). The baseline estimate ($\hat{b} = -1.11$) is stable across polynomial orders 5 through 8 (range: -1.11 to -2.11), bandwidths of $3–$5 (range: -1.11 to -2.24), and exclusion regions of $0.50–$1.50 (range: -0.74 to -1.11). The $2 bandwidth is too narrow to fit the polynomial reliably, and the $2.00 exclusion region changes sign, reflecting overfitting when the excluded zone is too wide relative to the bandwidth.

Placebo tests. Table 6 presents two sets of placebos. Panel A tests for bunching at round numbers ($5, $15, $20) that are not reporting thresholds. The $5 placebo shows positive $\hat{b} = 1.77$—round-number heaping that is mechanically different from the negative censoring

Table 6: Placebo Tests: Bunching at Non-Threshold Amounts

Test	Description	$\hat{b}$	SE($\hat{b}$)
Panel A: Round-Number Placebos			
	\$5	1.771	0.587
	\$15	0.461	0.238
	\$20	-0.441	0.275
Panel B: Prior-Year Threshold Placebos			
	2020 at \$10.42	-1.870	0.270
	2022 at \$10.73	316.657	2126.613
	2024 at \$11.29	-3.261	0.364

Notes: Panel A tests for bunching at round-dollar amounts that are not reporting thresholds. Panel B tests for bunching in year t at year $t-1$'s threshold (which no longer governs reporting). If bunching reflects strategic disclosure avoidance, it should track the current threshold, not prior thresholds or arbitrary round numbers.

discontinuity at the reporting threshold. Placebos at \$15 and \$20 are near zero. Panel B tests for bunching at prior-year thresholds in the current year. Results are mixed, with some prior-year thresholds still showing negative estimates—consistent with partial overlap between adjacent-year thresholds (\$0.30–\$0.50 apart).

Alternative explanations. Could the missing mass reflect something other than threshold-driven censoring? *Natural density decline:* The payment distribution may simply thin out at lower values, producing apparent missing mass that is not policy-driven. The polynomial counterfactual accounts for this smooth decline—the missing mass we measure is the deviation from the smooth fit, not the level of the density. *Round-number clustering:* If payments cluster at round dollars near the threshold (\$10, \$11, \$12), the polynomial might misattribute this clustering to the threshold. Our placebo tests show that round-number clustering produces *positive* bunching (excess mass at \$5), not the *negative* bunching we find at the reporting threshold. *Data truncation:* If CMS mechanically excludes all below-threshold payments, the density would drop to zero below the cutoff. Our data shows payments well below the threshold (down to \$2), confirming that the aggregate reporting rule does bring many sub-threshold payments into the database.

7. Discussion and Conclusion

The CMS Open Payments database is the most comprehensive public record of pharmaceutical industry-physician financial relationships. Researchers, journalists, and regulators rely on it to study conflicts of interest and their consequences for prescribing. We show that the database has a systematic blind spot: the per-transaction reporting threshold censors a substantial fraction of small payments from the public record.

The censoring is not symmetric across payment types. Food and beverage—the most frequent and arguably most consequential category for relationship maintenance (**?**)—shows the largest gap. This is the category where amounts are most elastic: a sales representative can trivially adjust a lunch order to keep the per-physician cost below \$11 or \$12. When the annual total from that manufacturer to that physician stays below \$100, the interaction vanishes entirely from the database. The public thus sees fewer of precisely the small, repeated contacts that the behavioral literature identifies as most influential.

Three implications follow. First, studies using Open Payments to measure the prevalence or intensity of industry-physician relationships likely understate the reach of pharmaceutical promotion at the bottom of the distribution. Effect estimates that rely on the *extensive margin*—whether a physician received any payment—are biased toward zero because some "untreated" physicians in fact received below-threshold meals that escaped the database. Second, the design flaw is fixable: eliminating the per-transaction exemption (reporting all payments regardless of size) or setting the threshold at \$1 would close the blind spot at minimal compliance cost. Third, similar threshold-driven censoring likely affects other disclosure regimes—lobbying expenditure reports, campaign finance databases, and financial advisor compensation disclosures—wherever a per-transaction floor determines what enters the public record.

The findings are subject to three important limitations. First, our sample covers approximately 160,000 records per year from the API, representing roughly 1 percent of each year's full dataset. The API's pagination order is not explicitly randomized, so the sample may not perfectly represent the population distribution. The consistency of results across four independently queried years—and the alignment of summary statistics with known aggregate properties of the database (e.g., the dominance of food and beverage payments)—mitigates but does not eliminate this concern. Future work should replicate these estimates using the complete bulk-download files. Second, the bunching estimator measures the density discontinuity at the threshold, not the total volume of unreported payments. Third, we cannot distinguish between mechanical censoring (payments simply not reported) and strategic avoidance (manufacturers deliberately sizing payments below the threshold). The cross-year tracking of the censoring point with the CPI-adjusted threshold is consistent with both mechanisms. Linking to Medicare Part D prescribing data—which we leave for future work—could help distinguish these channels. Many below-threshold payments are reported through the aggregate rule; our estimates capture only the marginal censoring at the per-transaction cutoff, not the full universe of unreported transfers.

When transparency depends on a bright line, the bright line becomes a blind spot. The Physician Payments Sunshine Act illuminated billions of dollars in previously hidden financial relationships between drug companies and the physicians who prescribe their products. But the per-transaction reporting floor ensures that the smallest, most frequent interactions—the ones that sales representatives can most easily size below the threshold—remain in the dark.

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: @ai1scl

References

Table 7: Standardized Disclosure Avoidance Estimates

Outcome	$\hat{b}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Pooled (all types)	-1.108	0.249	5.84	-1.108	0.249	Large negative
Panel B: Heterogeneous						
Food & Beverage only	-1.580	0.202	5.84	-1.580	0.202	Large negative
Non-food types	10.757	262.498	5.84	10.757	262.498	Large positive

Notes: **Country:** United States. **Research question:** Does pharmaceutical manufacturer disclosure avoidance at CMS Open Payments per-transaction reporting thresholds distort the distribution of physician payments? **Policy mechanism:** The Physician Payments Sunshine Act (ACA Section 6002) requires manufacturers to publicly report individual payments to physicians exceeding a CPI-indexed per-transaction minimum; payments below are exempt unless the annual aggregate exceeds a separate threshold, creating an incentive to keep individual payment amounts just below the per-transaction cutoff. **Outcome definition:** Normalized excess mass ($\hat{b}$) of the payment amount distribution just below the reporting threshold, measuring the share of payments shifted below the cutoff relative to the counterfactual density. **Treatment:** Binary: payment subject to per-transaction reporting threshold (CPI-adjusted, \$10.42 in 2018 to \$12.70 in 2024). **Data:** CMS Open Payments General Payment Data, 2018–2024, individual payment-level, payments in \$2–\$30 range. **Method:** Polynomial bunching estimator (order 7) with Poisson bootstrap standard errors (200 replications); exclusion region [threshold − \$1.50, threshold + \$0.50]. **Sample:** General (non-research, non-ownership) payments in the \$2–\$30 range; excludes payments above \$30 and below \$2. SDE = $\hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large (|SDE| > 0.15), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

A. Standardized Effect Sizes