The Fifty-Bed Cliff: How Medicare Payment Rules Shrink Rural Hospitals

APEP Autonomous Research* @ai1scl

March 30, 2026

Abstract

A hospital with 50 beds and one with 51 beds sit on opposite sides of a Medicare payment cliff: provider-based Rural Health Clinics at sub-50-bed hospitals receive uncapped cost-based reimbursement, while those at larger hospitals face per-visit caps. Using 50,398 hospital-year observations from CMS cost reports (2012–2023), I document dramatic bunching at this threshold. The count of hospitals at 50 beds is 5.7 times the count at 51 beds, implying that roughly 868 hospital-year observations represent facilities that would have operated above 50 beds absent the payment incentive. Placebo tests at non-regulatory thresholds (30, 40, 60, 75 beds) show no comparable distortion. The bunching predates the 2018 Bipartisan Budget Act that formalized the caps, suggesting hospitals had already internalized the cross-subsidy logic. These results reveal a capacity distortion where inpatient beds are sacrificed to preserve outpatient clinic revenue.

JEL Codes: I11, I18, H51

Keywords: hospital beds, bunching, Medicare, Rural Health Clinics, payment notch, capacity distortion

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 34m).

1. Introduction

In 2023, seventy-two American hospitals reported exactly 50 beds to CMS. Ten reported 51. This 7.2-to-1 ratio is not a coincidence—it is the visible signature of a payment rule that has quietly reshaped the American hospital landscape for over a decade.

Medicare's Rural Health Clinic (RHC) program creates a sharp discontinuity in hospital reimbursement at 50 beds. Provider-based RHCs affiliated with hospitals that maintain fewer than 50 beds receive uncapped cost-based reimbursement for outpatient clinic visits. When the host hospital crosses the 50-bed line, Medicare imposes per-visit payment caps on the affiliated RHC, substantially reducing revenue per encounter. The financial stakes are large: for a mid-sized rural hospital operating a provider-based RHC, the difference between uncapped and capped reimbursement can exceed several hundred thousand dollars annually.

This paper documents the resulting capacity distortion using the universe of non-Critical Access Hospital cost reports filed with CMS from fiscal years 2012 through 2023. I apply the bunching methodology of Kleven (2016) to the bed-count distribution and find an excess mass statistic of $\hat{b} = 2.25$, indicating that the density of hospitals at or just below the threshold is more than triple what the counterfactual distribution predicts. The implied reallocation is substantial: approximately 868 hospital-year observations in the data represent facilities that would have operated above 50 beds absent the payment incentive. These are hospitals that have chosen to constrain their inpatient capacity—forgoing beds, and by extension the patients who would fill them—in order to preserve uncapped outpatient reimbursement.

Several features of the setting make this distortion unusually clean to identify. First, the threshold is sharp and well-defined: the regulatory variable is the total licensed bed count reported on CMS Form 2552-10 (Worksheet S-3, Part I, Line 14), and the payment rule change occurs at exactly 50 beds. Second, I can construct convincing placebos. Bunching at round numbers (40, 60) that carry no regulatory significance is far smaller than at 50, and the distortion at 30 beds—where the distribution is smooth—is negative. Third, the 2018 Bipartisan Budget Act (BBA) provides a natural experiment in regulatory formalization: while the uncapped/capped distinction predated 2018, the BBA explicitly codified the per-visit caps. If hospitals were already responding to the underlying payment incentive, pre-BBA and post-BBA bunching should be comparable. That is precisely what I find ($\hat{b}_{pre} = 2.25$ versus $\hat{b}_{post} = 2.26$), suggesting that the distortion reflects longstanding hospital optimization, not a response to the 2018 rule change.

This paper contributes to three literatures. First, I add to the growing body of work using bunching methods to study behavioral responses to notches in tax and regulatory systems (Kleven and Waseem, 2013; Saez, 2010). While bunching has been extensively applied in public finance—at income tax kinks (Chetty et al., 2011), corporate size thresholds (Garicano et al., 2016), and charitable contribution limits—applications to healthcare regulation remain rare. The 50-bed threshold provides an unusually clean notch where the behavioral margin (bed count) is directly observable and precisely measured.

Second, the results speak to the literature on how Medicare payment design shapes hospital behavior (Duggan, 2000; Finkelstein, 2007). Bazzoli et al. (2008) documents bunching at the 100-bed threshold for the Disproportionate Share Hospital (DSH) program, but the 50-bed threshold operates through a different channel (outpatient clinic revenue rather than inpatient payment adjustments) and has not been previously studied. My finding that bunching was already fully developed before the 2018 BBA formalization suggests that hospitals respond to payment incentives embedded in regulatory rules well before those rules are explicitly codified—a form of anticipatory regulatory arbitrage.

Third, this paper contributes to the policy debate about rural hospital viability. The 2023 creation of the Rural Emergency Hospital (REH) designation, which is available only to hospitals with 50 or fewer beds, added a second incentive to stay below the threshold. Understanding the existing distortion is essential for predicting how the REH option will reshape rural healthcare capacity. If hospitals are already constraining bed counts for RHC payment reasons, the REH option compounds the incentive and may accelerate the downsizing of small rural hospitals—precisely the facilities that rural communities depend on for inpatient care.

The rest of the paper proceeds as follows. Section 2 describes the institutional background. Section 3 presents the data. Section 4 outlines the bunching methodology. Section 5 reports the main results and robustness checks. Section 6 discusses implications.

2. Institutional Background

The Rural Health Clinic Program. The Medicare RHC program, established by the Rural Health Clinic Services Act of 1977, subsidizes primary care in underserved rural areas by reimbursing qualifying clinics on a cost basis rather than through the standard fee schedule. There are two types of RHCs: independent clinics and provider-based clinics affiliated with a hospital. For provider-based RHCs, Medicare's reimbursement rules depend on the size of the affiliated hospital.

The 50-Bed Threshold. Provider-based RHCs affiliated with hospitals that have fewer than 50 beds receive cost-based reimbursement without a per-visit payment cap. Their Medicare reimbursement covers reasonable costs of furnishing RHC services, subject to standard prudent-buyer and cost-accounting rules but without a per-visit ceiling. When the host hospital reaches or exceeds 50 beds, Medicare imposes a per-visit payment cap (set at $92.51 for independent RHCs in 2024, with adjustments for provider-based clinics). This cap has been in place informally since the early 2000s, but was explicitly codified in Section 50901 of the Bipartisan Budget Act of 2018 (42 CFR 405.2462).

The 2018 Bipartisan Budget Act. The BBA formalized the distinction between small (<50 beds) and larger (≥50 beds) hospitals for RHC payment purposes. While hospitals appear to have already been responding to the informal threshold—as the pre-2018 bunching data demonstrate—the BBA removed any ambiguity about the rule's application.

The 2023 Rural Emergency Hospital Option. The Consolidated Appropriations Act of 2021 created the Rural Emergency Hospital (REH) designation, effective January 1, 2023. Hospitals with 50 or fewer beds may convert to REH status, receiving enhanced outpatient reimbursement plus an additional monthly facility payment of approximately $272,866. The REH option further reinforces the value of staying at or below 50 beds, potentially intensifying bunching at the threshold.

3. Data

I use the CMS Hospital Cost Report Information System (HCRIS), Form 2552-10, which covers all Medicare-participating hospitals in the United States. Each hospital files an annual cost report that includes, among thousands of line items, the total number of licensed beds (Worksheet S-3, Part I, Line 14, Column 2). I download individual fiscal year files from CMS for FY2012 through FY2023, yielding a panel of 50,398 non-CAH hospital-year observations covering 5,401 unique hospitals.

Sample Restrictions. I exclude Critical Access Hospitals (CAHs), which face their own 25-bed cap and receive distinct cost-based reimbursement. CAH status is identified from the CMS Certification Number (CCN), where digits 3–6 fall in the 1300–1399 range. I restrict to hospitals with positive bed counts and keep only reports with filed, settled, or reopened status. When a hospital files multiple reports for the same fiscal year, I retain the most recent report.

Period Definitions. I define two periods based on the BBA: *pre-BBA* (FY2012–2017) and *post-BBA* (FY2018–2023). The REH era begins in FY2023.

Table 1 reports summary statistics. The analysis window (20–100 beds) contains the majority of observations relevant to the 50-bed threshold. Notably, 10.2% of hospital-year

Table 1: Summary Statistics: Non-CAH Hospitals, 20–100 Beds

	Hospital-years	Unique hospitals	Mean beds	Median beds	SD beds	Share 46–50	Share 51–55
Pre-BBA (2012–2017)	11,500	2,469	55.0	50	22.3	11.9%	4.5%
Post-BBA (2018–2023)	11,690	2,424	54.4	49	22.1	12.8%	4.2%
Full sample (2012–2023)	23,190	2,887	54.7	50	22.2	12.4%	4.4%

Notes: Sample includes non-Critical Access Hospital (non-CAH) acute care hospitals with 20–100 beds from CMS HCRIS Form 2552-10 cost reports, fiscal years 2012–2023. The 46–50 and 51–55 columns show the share of hospital-year observations in each bin range. Under the null of no bunching, these shares should be approximately equal.

observations in the 20–100 bed range fall at 46–50 beds, compared to just 3.6% at 51–55 beds—a ratio of 2.8 that foreshadows the formal bunching analysis.

4. Methods

I follow the bunching estimation framework of Kleven (2016) and Kleven and Waseem (2013). The key idea is to estimate the counterfactual density of hospital bed counts—what the distribution would look like absent the payment notch at 50 beds—and compare it to the observed distribution to quantify the behavioral response.

Counterfactual Estimation. Let c_j denote the count of hospital-year observations in bed-count bin j, for $j \in \{20, 21, \ldots, 100\}$. I estimate the counterfactual density by fitting a flexible polynomial to the observed distribution, excluding bins in the manipulation region $[46, 55]$:

$$c_j = \sum_{k=0}^{p} \beta_k j^k + \sum_{i=j_L}^{j_U} \gamma_i \cdot \mathbb{I}[j = i] + \varepsilon_j \tag{1}$$

where $p = 7$ is the polynomial order and $[j_L, j_U] = [46, 55]$ is the exclusion window. The counterfactual density is the polynomial prediction $\hat{c}_j = \sum_{k=0}^{p} \hat{\beta}_k j^k$.

Excess Mass Statistic. The bunching estimate $\hat{b}$ is the ratio of excess mass in the bunching region (46–50) to the counterfactual density at the threshold:

$$\hat{b} = \frac{\sum_{j=46}^{50} (c_j - \hat{c}_j)}{\hat{c}_{50}} \tag{2}$$

An estimate of $\hat{b} = 2$ means that the observed density at the threshold is three times the counterfactual density—two extra “units” of density for every one that would naturally occur.

Inference. I compute standard errors by bootstrap resampling at the provider level. In each replication, I resample hospitals with replacement, reconstruct the bin counts, re-estimate the polynomial, and recompute $\hat{b}$. Provider-level resampling accounts for serial correlation in bed counts across years within the same hospital. The main estimates use 500 bootstrap replications; robustness checks use 200–300 replications for computational efficiency.

5. Results

The Fifty-Bed Cliff is Dramatic and Unambiguous. Table 2 presents the raw bed-count distribution from 40 to 60 beds. The pattern is striking: between 40 and 50 beds, counts are relatively stable (ranging from 228 to 1,125 per bin, with variation partly reflecting round-number heaping at multiples of 10). At the threshold, 776 hospital-year observations report exactly 50 beds. Immediately above, the count plummets: only 137 observations at 51 beds—a drop ratio of 5.7:1. The collapse persists through 55 beds (137 to 254 per bin), before recovering toward the round-number peak at 60 (899 observations).

The formal bunching estimates in Table 3 confirm the visual pattern. The pooled excess mass statistic is $\hat{b} = 2.253$ (SE = 0.404), statistically significant at the 1% level. This implies that the observed density in the bunching region is 225% above the counterfactual. The implied excess of 868 hospital-year observations means that, in a typical year, roughly 72 hospitals operate below the 50-bed threshold that would otherwise have 51 or more beds. For context, this is approximately 1.7% of all non-CAH hospitals in our sample—a meaningful share of the industry constrained by a single payment rule.

Bunching Predates the 2018 BBA. The pre-BBA (2012–2017) and post-BBA (2018–2023) estimates are nearly identical: $\hat{b}_{pre} = 2.250$ (SE = 0.418) versus $\hat{b}_{post} = 2.255$ (SE = 0.416), with a difference of $\Delta\hat{b} = 0.005$ (SE = 0.590, $t = 0.01$). This null result is informative: it rules out the hypothesis that the 2018 BBA created the distortion. Instead, the bunching appears to reflect hospital responses to the longstanding, informal payment distinction between sub-50 and 50+ bed facilities. The BBA merely codified an incentive that hospitals had already internalized.

Placebo Thresholds Confirm the Regulatory Channel. Table 4 Panel A reports bunching estimates at non-regulatory thresholds. No placebo produces bunching comparable to the 50-bed threshold ($\hat{b} = 2.253$). At 30 beds, the estimate is negative (-0.552); at 40 beds, it is 0.762; at 60 beds, it is 1.434 (some round-number heaping but well below the policy threshold); and at 75 beds, 0.427. The 50-bed estimate is roughly three times the largest placebo. The pattern is clear: the 50-bed bunching is regulatory in origin, not an

Table 2: Hospital Bed Count Distribution Near the 50-Bed Threshold

Bed count	Pre-BBA (2012–2017)	Post-BBA (2018–2023)	Total
40	506	619	1125
41	150	162	312
42	208	205	413
43	99	129	228
44	208	229	437
45	165	179	344
46	136	146	282
47	171	193	364
48	281	304	585
49	400	461	861
50	**379**	**397**	**776**
51	81	56	137
52	119	122	241
53	97	96	193
54	135	119	254
55	90	97	187
56	141	129	270
57	78	84	162
58	127	150	277
59	93	95	188
60	461	438	899
Drop ratio (50→51)	4.7:1	7.1:1	5.7:1

Notes: Hospital-year observations of non-CAH hospitals from HCRIS cost reports. Bold row marks the 50-bed threshold where provider-based RHC payment rules change. The horizontal line separates the uncapped reimbursement region (≤50 beds) from the capped region (>50 beds). The drop ratio reports the count at 50 beds divided by the count at 51 beds.

artifact of round-number reporting or general distributional features.

An additional test reinforces this conclusion. Across round numbers (20, 30, 40, 60, 70, 80), the average "heaping ratio" (count at the round number divided by the average of adjacent bins) is 3.48. At 50 beds, the comparable ratio is only 1.56. The lower heaping ratio at 50 is not evidence against bunching—it reflects the asymmetric distortion of the adjacent bins themselves. The bin at 49 (861 observations) is inflated by bunching, while the bin at 51 (137 observations) is depleted, compressing the ratio. This asymmetry is itself diagnostic of strategic behavior rather than mere rounding.

Table 3: Bunching Estimates at the 50-Bed Threshold

	Full sample (2012–2023)	Pre-BBA (2012–2017)	Post-BBA (2018–2023)
Excess mass ($\hat{b}$)	2.253***	2.250***	2.255***
	(0.404)	(0.418)	(0.416)
Excess hospital-years	867	415	453
	(83)	(42)	(46)
Missing mass above 50	761	339	422
Log density ratio	1.042	0.963	1.119
Polynomial order	7	7	7
Exclusion window	[46, 55]	[46, 55]	[46, 55]
Observations	50,398	25,646	24,752

Notes: Bunching estimates following Kleven (2016). The excess mass statistic $\hat{b}$ is the ratio of excess mass in the bunching region (46–50 beds) to the counterfactual density at the threshold, where the counterfactual is estimated from a degree-7 polynomial fit to the observed distribution excluding the manipulation region [46, 55]. Standard errors from 500 bootstrap replications (provider-level resampling) in parentheses. The log density ratio compares the count of hospital-years in the 5-bin window below (46–50) to the 5-bin window above (51–55) the threshold. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

Robustness to Specification Choices. Panel B of Table 4 shows that the bunching estimate is stable across polynomial orders 5 through 9 ($\hat{b}$ ranges from 2.25 to 2.87). Panel C varies the exclusion window: narrower windows yield larger estimates (3.32 for [47, 53]) because more of the distorted bins are excluded from the counterfactual, while wider windows yield smaller but still significant estimates (1.06 for [44, 57]). The density discontinuity test—a simple comparison of hospital counts in the 5-bin windows on either side of 50—yields a log ratio of 1.04 ($z = 28.5$), confirming that the density gap is massive and highly statistically significant.

Temporal Dynamics. Table 5 reports year-by-year bunching estimates. The distortion is present in every year from 2012 to 2023, with $\hat{b}$ ranging from 1.36 (2023) to 2.73 (2018). There is suggestive evidence that bunching peaked immediately after the BBA ($\hat{b} = 2.73$ in 2018, $\hat{b} = 2.62$ in 2019) before declining in 2022–2023. The 2023 decline may reflect early responses to the REH conversion option, which gives hospitals at or below 50 beds an alternative to maintaining inpatient services. However, with only one post-REH year, this interpretation is preliminary.

Table 4: Robustness: Placebo Thresholds, Polynomial Order, and Exclusion Windows

Panel	Specification	$\hat{b}$	SE
A. Placebo Thresholds			
	30 beds	-0.552	(0.400)
	40 beds	0.762	(0.327)
	50 beds (policy)	**2.253*****	**(0.368)**
	60 beds	1.434	(0.474)
	75 beds	0.427	(0.427)
B. Polynomial Order (threshold = 50)			
	Order 5	2.683***	(0.401)
	Order 6	2.548***	(0.404)
	Order 7 (baseline)	**2.253*****	**(0.350)**
	Order 8	2.541***	(0.452)
	Order 9	2.866***	(0.440)
C. Exclusion Window (threshold = 50, order 7)			
	[47,53]	3.322***	(0.427)
	[46,54]	2.603***	(0.385)
	[46,55] (baseline)	**2.253*****	**(0.412)**
	[45,56]	1.651***	(0.422)
	[44,57]	1.059***	(0.398)

Notes: Panel A reports bunching estimates at alternative thresholds with no regulatory significance (30, 40, 60, 75 beds) alongside the policy-relevant 50-bed threshold. Each placebo uses the same methodology: degree-7 polynomial with symmetric exclusion window. Panel B varies the polynomial order for the counterfactual density at the 50-bed threshold. Panel C varies the exclusion window. All estimates use the full 2012–2023 sample. Bootstrap SEs (300 replications) in parentheses.

6. Discussion and Conclusion

The fifty-bed cliff reveals a hidden cost of fragmented payment design. When Medicare reimburses outpatient clinics and inpatient services through separate channels with different size thresholds, hospitals optimize across margins in ways that policymakers may not intend. A hospital administrator facing the choice between adding a 51st bed (gaining inpatient capacity but losing uncapped RHC reimbursement) and staying at 50 beds (preserving the cross-subsidy) is making a rational financial calculation. The result—fewer inpatient beds in communities that may need them—is a cost borne by patients whose access to care depends on these facilities.

The economics of staying small. A rough calculation illustrates why hospitals find the cross-subsidy worth protecting. A typical provider-based RHC serves approximately 4,000–

Table 5: Year-by-Year Bunching Estimates at the 50-Bed Threshold

Fiscal year	$\hat{b}$	SE	Excess mass	Period
2012	2.435***	(0.678)	44	Pre-BBA
2013	2.444***	(0.530)	78	Pre-BBA
2014	2.359***	(0.544)	77	Pre-BBA
2015	2.492***	(0.516)	84	Pre-BBA
2016	1.806***	(0.496)	62	Pre-BBA
2017	2.078***	(0.465)	71	Pre-BBA
2018	2.732***	(0.504)	94	Post-BBA
2019	2.621***	(0.549)	91	Post-BBA
2020	2.315***	(0.481)	81	Post-BBA
2021	2.549***	(0.478)	83	Post-BBA
2022	1.871***	(0.460)	60	Post-BBA
2023	1.364***	(0.489)	44	Post-BBA

Notes: Year-by-year bunching estimates at the 50-bed threshold using degree-7 polynomial with exclusion window [46, 55]. The Bipartisan Budget Act of 2018 introduced explicit per-visit payment caps for provider-based RHCs at hospitals with 50 or more beds. The Rural Emergency Hospital conversion option became available in January 2023 for hospitals with 50 or fewer beds. Bootstrap SEs (200 replications) in parentheses.

8,000 Medicare visits per year. With the 2024 per-visit cap set at $92.51 for independent RHCs (provider-based caps are adjusted but comparable), and average cost-based reimbursement for uncapped RHCs substantially exceeding the cap, the annual revenue differential for a single hospital can plausibly range from $200,000 to $500,000. For a 50-bed rural hospital with thin operating margins, this revenue stream can represent the difference between solvency and closure. In aggregate, the 72 hospitals per year that appear to constrain their size collectively forgo an estimated 72–360 inpatient beds—capacity that communities served by these facilities cannot access.

Regulatory anticipation, not regulatory response. The most striking finding is not the bunching itself but its temporal invariance. The pre-BBA and post-BBA estimates are statistically indistinguishable ($\hat{b}_{pre} = 2.250$ versus $\hat{b}_{post} = 2.255$), implying that the 2018 BBA codified a behavioral response that hospitals had already fully adopted. This pattern—optimization in response to the *economic substance* of a rule, irrespective of its formal legal status—has broad implications for regulatory design. It suggests that CMS cannot assume payment thresholds are behaviorally inert simply because they have not been explicitly legislated. Any financial discontinuity embedded in Medicare's fee structure, however informal, should be presumed to generate behavioral responses proportional to the stakes.

Policy design implications. The sharp notch at 50 beds creates the distortion. A smooth transition—phasing the per-visit cap in gradually between, say, 45 and 55 beds—would reduce the incentive to cluster at the threshold while preserving the program's intent of supporting small rural hospitals. The 2023 REH conversion option, which adds a second layer of value to staying below 50 beds (approximately $3.3 million annually in facility payments), may compound the distortion. Policymakers should monitor whether the post-2023 bed distribution shows further compression below the threshold and consider decoupling REH eligibility from the same bed-count cutoff used for RHC payment rules.

Limitations. Several caveats warrant discussion. First, the analysis documents bunching in *licensed* bed counts, which reflect a hospital's regulatory ceiling rather than its day-to-day staffed capacity. Changing licensed bed counts is not costless—it typically requires state health department approval, updated certificates of need (in states that require them), and amendments to Medicare provider agreements—but it is an administrative action rather than a physical one. Hospitals can delicense beds without removing them from the building. The bunching therefore reflects strategic positioning of the regulatory variable, though the financial consequences of that positioning (RHC payment rules, REH eligibility) are real regardless of whether physical capacity changes.

Second, the analysis treats all non-CAH hospitals equally, but the 50-bed notch matters only for hospitals that operate provider-based RHCs. If hospitals without RHCs also bunch at 50 beds, the regulatory mechanism I propose would be incomplete. Linking the HCRIS data to CMS RHC enrollment files to compare bunching among RHC-affiliated versus non-affiliated hospitals is a natural extension that would strengthen the causal interpretation.

Third, the panel runs only through 2023, providing limited power to evaluate the REH effect. Fourth, while I exclude CAH hospitals, other Medicare designations (Sole Community Hospital, Medicare Dependent Hospital) carry their own size-related rules. I am not aware of any such rule with a discontinuity at exactly 50 beds, but the possibility of confounding regulatory thresholds cannot be fully excluded without a comprehensive audit of all bed-count-dependent Medicare provisions.

References

Bazzoli, Gloria J, Hsueh-Fen Chen, Man Zhao, and Richard C Lindrooth, "Hospital Financial Condition and the Quality of Patient Care," *Health Economics*, 2008, *17* (8), 977–995.

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.

Duggan, Mark G, "Hospital Ownership and Public Medical Spending," *Quarterly Journal of Economics*, 2000, *115* (4), 1343–1373.

Finkelstein, Amy, "The Effect of Public Insurance on Morally Hazardous Behavior: Medicare and Hospital Stays," *American Economic Review*, 2007, *97* (5), 1912–1933.

Garicano, Luis, Claire Lelarge, and John Van Reenen, "Firm Size Distortions and the Productivity Distribution: Evidence from France," *American Economic Review*, 2016, *106* (11), 3439–3479.

Kleven, Henrik J, "Bunching," *Annual Review of Economics*, 2016, *8*, 435–464.

_ **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.

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

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled (2012–2023)						
Bunching intensity ($\hat{b}$)	2.253	0.404	184.220	4.709	0.844	Large positive
Log density ratio	1.042	0.037	1.000	1.042	0.037	Large positive
Panel B: Heterogeneous (Pre/Post BBA 2018)						
Pre-BBA $\hat{b}$ (2012–2017)	2.250	0.418	87.268	4.754	0.883	Large positive
Post-BBA $\hat{b}$ (2018–2023)	2.255	0.416	97.983	4.619	0.852	Large positive

Notes: **Country:** United States. **Research question:** Does the Medicare Rural Health Clinic payment threshold at 50 hospital beds distort hospital capacity decisions, causing hospitals to maintain fewer beds than they otherwise would? **Policy mechanism:** Provider-based Rural Health Clinics affiliated with hospitals below 50 beds receive uncapped cost-based Medicare reimbursement, while those at 50 or more beds face per-visit payment caps, creating a financial incentive to constrain inpatient bed capacity to preserve outpatient clinic revenue. **Outcome definition:** Bunching statistic $\hat{b}$ measuring excess mass of hospitals at or just below the 50-bed threshold relative to the counterfactual density estimated from a polynomial fit to the distribution excluding the manipulation region. **Treatment:** Binary—hospital bed count falling below versus at-or-above the 50-bed threshold. **Data:** CMS Hospital Cost Report Information System (HCRIS) Form 2552-10, fiscal years 2012–2023, non-Critical Access Hospital acute care hospitals with 20–100 beds (N = 50,398 hospital-year observations, 5,401 unique hospitals). **Method:** Bunching estimation following Kleven (2016) with degree-7 polynomial counterfactual, exclusion window [46, 55], provider-level bootstrap standard errors (500 replications). **Sample:** Non-CAH hospitals reporting to CMS with positive bed counts; restricted to 20–100 beds for polynomial estimation, with results robust to wider windows. SDE = $\hat{\beta}$/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

Acknowledgements

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

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

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