Tariff Cliffs: Hidden Notches in the UK's Solar Subsidy

APEP Autonomous Research* @SocialCatalystLab

March 31, 2026

Abstract

Britain's Feed-in Tariff paid a single generation rate based on total installed capacity, creating *tariff cliffs*—hidden notches where crossing a capacity-band boundary lowered the rate on all kilowatt-hours, not just the marginal one. Using 860,470 solar PV installations, I document extreme bunching below the 4 kW threshold: during 2010–2015, only 0.6% of domestic-scale installations exceeded 4 kW, and the raw bunching ratio reached 2,230:1. The February 2016 band merger eliminated the 4 kW cliff. Bunching collapsed: the share above 4 kW rose to 12%, the ratio fell to 20.6:1 by 2017. Engineering mass points at nearby capacities did not collapse, confirming the policy origin. At the unchanged 10 kW threshold, Kleven-Waseem bunching persisted at 56.5 (FIT period) and 36.3 (post-merger). Average-rate subsidy design can create first-order sizing distortions invisible in the tariff schedule itself.

JEL Codes: H23, Q42, Q48, Q58

Keywords: bunching, solar energy, feed-in tariffs, renewable energy policy, tariff design, hidden notches

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 47m).

1. Introduction

Feed-in tariffs are designed to encourage renewable energy by guaranteeing a price per kilowatt-hour. When tariff rates differ across capacity bands, the government intends a graduated schedule: larger systems receive modestly lower rates, reflecting declining per-unit costs. But if the tariff rate applies to the *entire* installation based on total capacity—rather than to each kilowatt-hour at its own marginal band—the graduated schedule conceals a series of discrete revenue cliffs. Crossing a band boundary does not just lower the rate on the next kilowatt; it lowers the rate on every kilowatt the system produces. These *tariff cliffs* are hidden notches: the subsidy schedule looks smooth, but the revenue function is not.

This paper documents extreme bunching at the 4 kW tariff cliff in the UK Feed-in Tariff (FIT). The UK FIT, operational from April 2010 to March 2019, assigned solar PV installations to one of several capacity bands—≤4 kW (highest rate), 4–10 kW, 10–50 kW, and >50 kW—with a single generation tariff rate applied to the installation's Declared Net Capacity (DNC). A homeowner with a 3.99 kW system received the ≤4 kW rate on all generation; a homeowner with a 4.01 kW system received the lower 4–10 kW rate on all generation. At the initial FIT rates (2010–2011), this difference cost approximately £190 per year in lost generation revenue—roughly £2,800 in net present value over the 20-year tariff guarantee. The cost of the additional 0.02 kW of panels: negligible. The tariff cliff was a one-sided bet, and the market responded accordingly.

During the FIT tariff-band regime (2010–2015), only 0.6% of installations in the domestic-scale window (3.5–5.0 kW DNC) exceeded 4.0 kW. The raw bunching ratio at 4.0 kW—the count in the threshold bin divided by the average count per 0.1 kW bin just above—reached 2,230:1 in 2012 and remained above 700:1 throughout 2011–2015. Above the cliff, the capacity distribution was virtually empty. Fewer than 100 installations per year were commissioned with DNC between 4.1 and 4.5 kW, in a market that was installing 40,000–70,000 systems annually in the same capacity range below the threshold.

The February 2016 FIT reform provides the identification. The government merged the ≤4 kW and 4–10 kW tariff bands into a single 0–10 kW band at a unified rate, eliminating the 4 kW cliff while leaving the 10 kW and 50 kW thresholds intact. Bunching at 4 kW collapsed: the raw ratio fell from 792:1 (2015) to 60.7:1 (2016, a transition year), to 20.6:1 (2017), to 2.2:1 (2019). The share of installations above 4 kW in the domestic window rose from 0.6% to 12%. The right tail came back.

Three pieces of evidence support the policy interpretation. First, engineering mass points at nearby capacities that are not tariff boundaries—3.68 kW (the single-phase connection standard), 3.92 kW, 3.99 kW (common module configurations)—did not collapse after the reform. Their shares increased, consistent with the removal of the cliff allowing more natural sizing variation. Second, at the unchanged 10 kW threshold, where the tariff cliff persisted throughout, Kleven-Waseem bunching estimates were 56.5 (pooled FIT period) and 36.3 (pooled post-merger)—materially positive in both regimes. Third, raw ratios at non-policy round numbers (3, 5, 6, 8 kW) showed stable values of 2–13 across the reform—no threshold effect, no reform response.

This paper contributes to the bunching literature (Saez, 2010; Kleven and Waseem, 2013) by documenting an extreme response that arises from a design feature—average-rate tariff banding—that policymakers may not have intended as a notch. The contribution is diagnostic: the three conditions that generate extreme bunching in the German solar market (Klimsa et al., 2024)—professional intermediation, modular technology, and disproportionate stakes—are present in the UK, and the average-rate structure amplifies them by making the effective notch size proportional to total system output rather than marginal output. The temporal on/off experiment at 4 kW, with 10 kW as an unchanged control, provides cleaner identification than cross-sectional bunching alone. Srivastav (2024) documents bunching at the UK FIT's 5 MW ceiling among utility-scale developers; this paper extends the bunching lens to the three domestic capacity thresholds where 95% of UK FIT installations occur.

Related Literature

The bunching literature, following Saez (2010) and Kleven and Waseem (2013), has developed tools for estimating behavioral responses at tax and regulatory thresholds. Key contributions include the distinction between kinks (marginal rate changes) and notches (discrete jumps), with notches predicted to generate larger responses (Kleven, 2016). The present paper identifies a hidden notch—one that appears as a kink in the tariff schedule but functions as a notch in the revenue function because the rate applies to total, not marginal, output.

Garicano et al. (2016) document bunching at France's 50-employee thresholds, where adjustment costs moderate the response. Best and Kleven (2018) find large distortions at UK stamp duty notches, exploiting temporal variation from temporary tax changes. Chetty et al. (2011) show that intermediation and adjustment costs shape the magnitude of threshold responses. The UK solar setting combines all three features—professional installers, near-zero adjustment costs, and disproportionate financial stakes—producing extreme bunching.

In the renewable energy literature, Klimsa et al. (2024) document bunching at Germany's 10 kWp solar surcharge threshold and estimate 69 MWp/year of foregone capacity from abolition. The German threshold is a pure notch (surcharge exemption below 10 kWp); the UK thresholds are hidden notches (average-rate tariff bands). Srivastav (2024) studies the UK FIT's 5 MW ceiling as a risk-reduction instrument, finding that the FIT induced at least 2.3 GW of additional solar capacity, but focuses on utility-scale developers rather than domestic installations.

2. Institutional Background

2.1 The UK Feed-in Tariff

The UK Feed-in Tariff, established by the Energy Act 2008 and operational from April 1, 2010 to March 31, 2019, guaranteed fixed payments per kilowatt-hour of electricity generated by small-scale renewable installations. For solar PV, the tariff consisted of two components: a *generation tariff* paid for every kWh generated (regardless of whether it was consumed on-site or exported) and a smaller *export tariff* paid for kWh exported to the grid. The generation tariff was the dominant revenue stream.

Tariff bands and the average-rate structure. Installations were assigned to capacity bands based on their Declared Net Capacity (DNC)—the maximum export capacity, determined by the inverter rating rather than panel capacity. The initial (2010–2011) generation tariff rates for solar PV were: ≤4 kW retrofit: 41.3 p/kWh; ≤4 kW new build: 36.1 p/kWh; 4–10 kW: 36.1 p/kWh; 10–50 kW: 31.4 p/kWh; 50–100 kW: 29.3 p/kWh. Critically, a single rate applied to the installation's entire generation based on its band assignment. A system with DNC of 3.99 kW received the ≤4 kW rate on all kWh; a system with DNC of 4.01 kW received the lower 4–10 kW rate on all kWh.

Degression. From 2012, tariff rates were reduced quarterly based on deployment volumes in each band ("degression"). Rates also varied by Energy Performance Certificate (EPC) rating (High, Medium, Low). By 2015, the ≤4 kW rate had fallen to approximately 12–13 p/kWh, narrowing the absolute differential at the 4 kW threshold but not eliminating it.

The February 2016 band merger. Following a comprehensive review, the government implemented major changes effective February 8, 2016 (with a scheme pause from January 15 to February 7). The ≤4 kW and 4–10 kW bands were merged into a single 0–10 kW band at a unified rate of approximately 4 p/kWh. This eliminated the 4 kW tariff cliff. The 10 kW, 50 kW, and 100 kW thresholds remained: the 10–50 kW band received approximately 2–3 p/kWh, creating a continuing tariff cliff at 10 kW.

2.2 The Installer Channel

Residential solar installations in the UK are designed and installed by MCS-certified professionals. The installer, not the homeowner, selects the number of panels and the inverter size that determine DNC. This intermediation is important because (a) installers face the tariff schedule repeatedly across dozens of installations and learn to optimize around thresholds; (b) offering a just-below-threshold system is a competitive advantage, since the homeowner sees higher revenue per kW; and (c) the technology is modular—removing or adding a single 300–400 Wp panel changes DNC by less than 0.5 kW at near-zero cost.

Table 1: Summary Statistics: Ofgem FIT Solar PV Installations

	FIT Bands (Apr 2010–Jan 2016)	Post-Merger (Feb 2016–Mar 2019)
Total installations	765,176	90,751
Domestic installations	736,284	84,761
Median DNC (kW)	3.30	3.51
Mean DNC (kW)	5.33	7.55
DNC distribution		
$\leq$ 4 kW	727,067 (95.0%)	76,571 (84.4%)
4–10 kW	16,328 (2.1%)	9,132 (10.1%)
10–50 kW	18,654 (2.4%)	4,515 (5.0%)
$>$ 50 kW	3,127 (0.4%)	533 (0.6%)
Missing tail: share above 4.0 kW (in [3.5, 5.0] window)		
Above 4.0 kW	0.6%	12.0%

Notes: Data from Ofgem FIT Installation Report (September 2025). DNC = Declared Net Capacity. FIT bands: April 2010 to January 2016. Post-merger: February 2016 to March 2019. Missing tail = share of installations in [3.5, 5.0] kW with DNC above 4.0 kW.

3. Data

I use the Ofgem Feed-in Tariff Installation Report (September 30, 2025 release), which provides installation-level records for every FIT-accredited installation in Great Britain. The dataset contains 860,470 solar PV installations commissioned between 2010 and 2019.

Key variables. The primary variable is Declared Net Capacity (DNC) in kW, the capacity measure that determines tariff band assignment. Commissioning date provides the temporal dimension. Tariff codes link each installation to its assigned band and rate. Postcode and LSOA enable geographic analysis.

Sample construction. The analysis uses all 860,470 solar PV installations. For the 4 kW analysis, the main sample restricts to installations with DNC between 1 and 7 kW. For the 10 kW analysis, the window is 5–15 kW. The FIT bands period runs from April 2010 to January 14, 2016 (765,176 installations); the post-merger period from February 8, 2016 to March 31, 2019 (90,751 installations).

4. Empirical Strategy

4.1 The Tariff Cliff as a Hidden Notch

The UK FIT's average-rate tariff structure creates a discrete revenue jump at each band boundary. Consider a homeowner choosing between a 3.99 kW system (in the ≤4 kW band) and a 4.01 kW system (in the 4–10 kW band). At initial rates:

- Revenue at 3.99 kW: $3.99 \times 950 \times 0.413 =$ £1,565/year
- Revenue at 4.01 kW: $4.01 \times 950 \times 0.361 =$ £1,375/year
- Annual revenue loss from crossing: £190
- NPV over 20 years at 3%: approximately £2,830
- Cost of the additional 0.02 kW ($\approx 1/15$ of a panel): negligible

The revenue function is discontinuous at 4 kW despite the tariff schedule appearing graduated. The effective notch size—the net present value of the revenue loss—is proportional to total system output, not just the marginal kilowatt. This distinguishes the UK's hidden notch from the German surcharge exemption studied by Klimsa et al. (2024), where the notch is an explicit surcharge.

4.2 Identification: The 2016 Band Merger

The February 2016 reform eliminated the 4 kW cliff by merging the ≤4 kW and 4–10 kW bands. This creates a natural experiment:

1. *Prediction 1:* Bunching at 4 kW should collapse after February 2016.
2. *Prediction 2:* Bunching at 10 kW should persist (threshold unchanged).
3. *Prediction 3:* Engineering mass points at non-policy capacities (3.68, 3.92, 3.99 kW) should not collapse.

All three predictions are confirmed in the data.

4.3 Estimation

At 4 kW, the domestic capacity distribution is heavily concentrated in a narrow range, with large engineering mass points from standard module configurations. The smooth counterfactual polynomial required by the standard Kleven-Waseem estimator is not credible in this region. Accordingly, the primary 4 kW estimands are: (a) the raw bunching ratio—the count at DNC = 4.0 divided by the average count per 0.1 kW bin in $[4.1, 4.5)$; (b) the missing-tail share—the fraction of installations in $[3.5, 5.0]$ with DNC above 4.0 kW; and (c) the discrete change in these objects at the 2016 reform.

At 10 kW, where the distribution is sparser and smoother, I apply the standard Kleven-Waseem bunching estimator following Kleven and Waseem (2013). A seventh-degree polynomial is fit to the capacity distribution outside an exclusion window of $[9.5, 10.5)$ kW, providing a counterfactual density. The bunching ratio $\hat{b} = \hat{B}/\hat{f}_0$ measures excess mass relative to counterfactual density at the threshold. Bootstrap standard errors use 500 replications.

5. Results

5.1 The 4 kW Tariff Cliff

Table 2 presents the annual evidence. During the FIT bands period (2010–2015), the raw bunching ratio at 4 kW rose from 21.6:1 in the scheme's first year to 2,230:1 in 2012, settling at 700–1,500:1 for the remainder of the period. The share of installations at exactly 4.0 kW within the $[4.0, 4.5)$ window exceeded 99% in every year from 2011 to 2015. The capacity distribution above 4 kW was virtually empty: in the 3.5–5.0 kW window, only 0.24% of installations exceeded 4.0 kW in 2012.

The pre-reform trend in the raw ratio was declining—from 2,230:1 in 2012 to 792:1 in 2015—consistent with tariff degression narrowing the notch. But the decline was gradual; the 2016 collapse was abrupt: from 792:1 to 60.7:1 in a single year, coinciding precisely with the band merger. No degression event during 2011–2015 produced a comparable discontinuity.

After the February 2016 band merger, the cliff collapsed. The transition year (2016) shows

Table 2: The 4 kW Tariff Cliff: Annual Evidence

Year	At 4.0 kW (count)	Avg above (per 0.1 bin)	Raw ratio	Share at 4.0 (%)	Share above 4.0 (%)
2010	207	10	21.6	81.2	2.4
2011	13,518	20	669.2	99.3	0.3
2012	20,074	9	2,230	99.8	0.2
2013	17,504	12	1,459	99.7	0.6
2014	23,269	17	1,402	99.6	0.6
2015	21,690	27	791.6	99.4	0.9
2016	3,851	63	60.7	92.4	4.9
2017	2,044	99	20.6	80.4	11.4
2018	1,973	167	11.8	70.2	13.1
2019	386	175	2.2	30.6	16.8

Notes: At 4.0 kW counts installations with DNC in $[4.0, 4.1)$ kW. Avg above is the average count per 0.1 kW bin in $[4.1, 4.5)$. Raw ratio = (At 4.0) / (Avg above). Share at 4.0 = fraction of installations in $[4.0, 4.5)$ that fall in the $[4.0, 4.1)$ bin. Share above 4.0 = fraction of installations in $[3.5, 5.0]$ with DNC strictly above 4.0 kW. Horizontal line separates the FIT bands period (pre-reform) from the post-merger period (post-reform). The February 2016 band merger eliminated the 4 kW tariff threshold by merging the $\leq$4 kW and 4–10 kW bands into a single 0–10 kW band. 2016 is a transition year (merger effective February 8).

a ratio of 60.7:1, reflecting the pre-reform installations commissioned in January. By 2017, the ratio had fallen to 20.6:1; by 2019, it was 2.2:1—within the range observed at non-policy round numbers. The share of installations above 4.0 kW rose from 0.6% (FIT period pooled) to 12% (post-merger pooled), reaching 16.8% by 2019.

5.2 Unchanged Control: 10 kW

The 10 kW threshold, which was not affected by the 2016 band merger, provides an internal control. Table 3 compares the annual bunching patterns at 4 kW and 10 kW. Pooled Kleven-Waseem estimates at 10 kW are 56.5 during the FIT period and 36.3 post-merger; annual estimates range from 41 to 70 during 2011–2015 and 25 to 51 during 2016–2019 (excluding 2010, where the thin ramp-up sample yields an imprecise 12.8). The attenuation reflects tariff degression reducing the economic significance of the 10 kW cliff, not the elimination of the threshold. The key contrast: the 4 kW raw ratio collapses by two orders of magnitude at

Table 3: Annual Bunching: 4 kW (Raw Ratio) vs. 10 kW (Kleven-Waseem)

Year	4 kW raw ratio	10 kW $\hat{b}$	(SE)
2010	21.6	12.8	
2011	669.2	65.1	(4.2)
2012	2,230	69.7	(5.3)
2013	1,459	61.0	(8.0)
2014	1,402	64.8	(7.4)
2015	791.6	41.3	(3.5)
2016	60.7	34.3	(5.9)
2017	20.6	37.4	(9.3)
2018	11.8	51.4	(21.4)
2019	2.2	25.0	(4.7)

Notes: The 4 kW raw ratio = count at DNC $\in [4.0, 4.1)$ / average count per 0.1 kW bin in $[4.1, 4.5)$. The 10 kW $\hat{b}$ is the Kleven-Waseem bunching ratio (polynomial degree 7, exclusion window $[9.5, 10.5)$ kW, estimation window $[5, 15)$ kW). Bootstrap standard errors (500 replications) in parentheses. The February 2016 reform eliminated the 4 kW threshold (band merger) but left the 10 kW threshold intact. The 4 kW ratio collapses post-reform; the 10 kW ratio persists.

the 2016 reform, while the 10 kW bunching ratio changes by less than a factor of two.

5.3 Within-Threshold Placebos

If the collapse of bunching at 4.0 kW were driven by a change in recording practices, rounding conventions, or market composition, it would also affect mass points at nearby capacities. It does not. The shares of installations at 3.68 kW (the single-phase connection standard) and 3.99 kW (a common module configuration) did not collapse after the reform—3.68 kW rose from 13.7% to 27.4% and 3.99 kW from 3.7% to 8.1%. These engineering mass points are not policy-dependent and serve as placebos: only the 4.0 kW mass point—the tariff boundary—responds to the policy change. Raw ratios at non-policy round numbers (3, 5, 6, 8 kW) show stable values of 2–13 across regimes, with no response to the 2016 reform (Table 4, Panel B).

6. Mechanism Evidence

6.1 DNC versus Installed Capacity

The Ofgem data include both Declared Net Capacity (the tariff-determining variable) and installed capacity (the physical panel capacity). During the FIT period, 21.8% of installations have installed capacity exceeding DNC. Among installations with DNC exactly 4.0 kW, the median installed capacity is also 4.0 kW, but the 95th percentile is 4.0 kW and the maximum reaches 8.0 kW. Some installers placed up to 8 kW of panels on the roof but registered a DNC of 4.0 kW by installing a smaller inverter, ensuring the system remained in the higher tariff band. This inverter-limiting strategy is a second adjustment margin beyond panel removal, underscoring the sophistication of the installer response.

6.2 The Nature of the Distortion

The tariff cliff did not prevent solar adoption—the UK installed over 736,000 domestic PV systems during the FIT bands period. Rather, it constrained system sizes. Systems that might have been designed at 4.5, 5, or 6 kW were instead capped at 4.0 kW to capture the higher tariff rate. The sizing distortion is a first-order behavioral response—among the most extreme in the bunching literature. A back-of-the-envelope calculation, using the post-merger distribution as the counterfactual, implies roughly 17–19 MW of foregone solar capacity during the FIT bands period. If the 12% above-4 kW share observed post-merger had prevailed during 2010–2015, roughly 35,000 additional installations would have been sized above 4 kW at an average excess of 0.5 kW—enough to power approximately 5,000 households. The capacity loss is modest compared to the German case documented by Klimsa et al. (2024), because UK residential systems were already small, but the behavioral response is more extreme. A bunching ratio of 2,230:1 at a tariff kink exceeds any estimate in the published bunching literature. The explanation lies in the hidden-notch structure: the effective notch size is proportional to total generation, not marginal generation, and the adjustment cost is near zero.

7. Robustness

Table 4, Panel A, reports a specification family for the Kleven-Waseem bunching estimate at 10 kW across polynomial degrees 6–8 and exclusion windows $[9.5, 10.5)$, $[9.0, 11.0)$, and $[8.5, 11.5)$ kW. The bunching ratio ranges from 7.6 to 84.2, with the narrow-window estimates (48–57) most stable. The wide-window, high-degree specification (degree 8, $[8.5, 11.5)$) produces an outlier at 7.6; the remaining eight specifications range from 36.8 to 84.2. The substantive conclusion—persistent, economically meaningful bunching at 10 kW—is robust.

The 4 kW results are inherently robust to estimator choice because they do not rely on a polynomial counterfactual. The raw ratio collapses from 954:1 (FIT pooled) to 14:1 (post-merger pooled) regardless of bin width, window definition, or comparison group.

8. Discussion

8.1 Average-Rate Design as a Hidden Notch

The standard bunching literature distinguishes kinks (marginal rate changes) from notches (discrete jumps), predicting that notches generate larger responses because they create dominated regions (Kleven, 2016). The UK FIT tariff schedule appears to create kinks—the rate changes at band boundaries, not the total payment. But because the rate applies to total capacity, not marginal capacity, the revenue function has discrete downward jumps at each boundary. The effective notch size is the rate differential multiplied by total output, which for a 4 kW system at initial rates exceeds £2,000 in NPV. This hidden-notch structure may be common in regulatory designs that assign entities to categories based on aggregate attributes: income-tax brackets that apply the top rate to all income (common in developing countries), utility tariffs based on total consumption, or subsidy tiers based on firm size.

8.2 Cross-Country Comparison

The German solar surcharge exemption at 10 kWp, documented by Klimsa et al. (2024), is an explicit notch: systems below 10 kWp are exempt from a surcharge that applies to all

Table 4: Robustness and Placebo Tests

Panel A: Specification Family at 10 kW (FIT Period)			
Degree	Window	$\hat{b}$	Excess Mass
6	$[9.5, 10.5)$ kW	56.8	6,157
6	$[9, 11)$ kW	84.2	6,848
6	$[8.5, 11.5)$ kW	37.3	5,216
7	$[9.5, 10.5)$ kW	**56.5**	**6,157**
7	$[9, 11)$ kW	82.7	6,848
7	$[8.5, 11.5)$ kW	36.8	5,216
8	$[9.5, 10.5)$ kW	48.3	6,014
8	$[9, 11)$ kW	63.1	6,520
8	$[8.5, 11.5)$ kW	7.6	2,169
Range		7.6–84.2	
Panel B: Placebo at Non-Policy Thresholds (Raw Ratios)			
Threshold		FIT Bands	Post-Merger
3 kW		2.1	3.0
5 kW		5.9	6.0
6 kW		8.1	12.9
8 kW		7.1	4.8
4 kW (policy)		954.0	14.1

Notes: Panel A reports the Kleven-Waseem bunching ratio at 10 kW across a pre-specified estimator family (3 polynomial degrees × 3 exclusion windows). Baseline specification (degree 7, $[9.5, 10.5)$ kW) in bold. Panel B reports raw ratios (count at threshold / average count per 0.1 kW bin in the 5 bins above) at round-number capacities that are not FIT tariff boundaries, compared with the policy threshold at 4 kW. Non-policy thresholds show stable ratios of 2–13 across regimes; only the 4 kW policy threshold shows a 68-fold collapse.

self-consumed electricity above the threshold. The estimated bunching ratio is approximately 87 under the notch and 13 under the preceding kink. The UK's hidden notch at 4 kW produces a raw ratio of 954:1 (pooled) and up to 2,230:1 in peak years—an order of magnitude larger. The difference reflects the hidden-notch structure (the entire tariff income is at stake, not just a surcharge) and the dominance of small domestic systems in the UK market (median DNC 3.5 kW), which concentrates the mass near the threshold.

8.3 Limitations

Three limitations merit acknowledgment. First, the raw bunching ratio at 4 kW, while dramatic, is not a structural elasticity estimate. The domestic capacity distribution is too concentrated near the threshold for the standard polynomial counterfactual to be credible, so I report raw ratios and missing-tail shares rather than Kleven-Waseem estimates. Second, I lack a tariff-rate crosswalk to precisely quantify the notch size at each threshold and in each degression period. The qualitative pattern—collapse at the eliminated threshold, persistence at the unchanged threshold—does not depend on the exact rate differential, but a crosswalk would sharpen the welfare calculation. Third, the 50 kW threshold confounds a tariff cliff with an administrative-regime change (MCS-FIT versus ROO-FIT accreditation), so it is excluded from the main analysis.

9. Conclusion

Average-rate tariff bands look smooth. They are not. When a subsidy assigns a single rate to the entire output of a technology based on total capacity, crossing a band boundary costs not one kilowatt-hour of subsidy but all of them. The UK's Feed-in Tariff created tariff cliffs at capacity boundaries that the tariff schedule itself did not reveal. The 4 kW cliff produced bunching ratios among the most extreme ever documented—and it vanished, on schedule, when the government merged the bands.

The diagnostic is portable. Any subsidy, tax, or regulatory schedule that assigns a single rate or status to the entire entity based on an aggregate attribute—total capacity, total income, total employment, total emissions—creates a hidden notch at the category boundary. Whether the notch distorts behavior depends on the same three conditions identified in the German solar market: sophisticated intermediaries, modular adjustment, and disproportionate stakes. The UK's experience shows that even a tariff schedule designed to *encourage* renewable deployment can constrain it, if the threshold design is not taken seriously.

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: @SocialCatalystLab

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

References

Best, Michael Carlos and Henrik Jacobsen Kleven, "Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK," *The Review of Economic Studies*, 2018, *85* (1), 157–193.

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," *The Quarterly Journal of Economics*, 2011, *126* (2), 749–804.

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 and Mazhar Waseem, "Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan," *The Quarterly Journal of Economics*, 2013, *128* (2), 669–723.

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

Klimsa, Drahomir, Mario Rieger, and Robert Ullmann, "How (not) to tax sunshine: Bunching around tax-exempt thresholds for rooftop photovoltaic systems," 2024. SSRN Working Paper 5022002.

Saez, Emmanuel, "Do taxpayers bunch at kink points?," *American Economic Journal: Economic Policy*, 2010, *2* (3), 180–212.

Srivastav, Sugandha, "Bringing breakthrough technologies to market: Risk reduction for solar power," 2024. Working Paper, University of Oxford.

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Missing-tail share (4 kW)	0.107	—	0.063	1.712	—	Large
Raw ratio collapse (4 kW)	−940	—	738	−1.27	—	Large
KW bunching diff (10 kW)	15.4	—	19.2	0.80	—	Large
Panel B: Heterogeneous						
Missing-tail (domestic)	0.107	—	0.063	1.712	—	Large
Missing-tail (non-domestic)	—	—	—	—	—	—
KW at 10 kW (pre-2013)	49.2	—	—	—	—	—

Notes: SDE = $\hat{\beta}/\text{SD}(Y)$. Classification: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005). Classification refers to magnitude, not statistical significance. **Country:** United Kingdom. **Research question:** Do average-rate FIT tariff bands create hidden notches that distort solar PV system sizing? **Policy mechanism:** FIT capacity-band tariff rates create discrete revenue cliffs at band boundaries. **Outcome definition:** Share of installations with DNC above 4.0 kW in a local window; Kleven-Waseem bunching ratio at 10 kW. **Treatment:** February 2016 band merger eliminating 4 kW threshold. **Data:** Ofgem FIT Installation Report, 860,470 solar PV installations (2010–2019). **Method:** Bunching estimation (Kleven-Waseem at 10 kW); missing-tail share comparison at 4 kW. **Sample:** 765,176 FIT-band installations; 90,751 post-merger installations.

A. Standardized Effect Sizes