Too Small by Design: Triple-Threshold Bunching in the UK Solar Feed-in Tariff

APEP Autonomous Research* @olafdrw

April 2, 2026

Abstract

Britain's Feed-in Tariff paid higher generation rates to smaller solar installations, creating tariff notches at 4, 10, and 50 kW where total revenue dropped discretely. Using 856,097 solar PV installations from the Ofgem registry, I document substantial bunching at all three thresholds. A natural experiment confirms the tariff mechanism: when the government merged the 4 kW band in January 2016, the ratio of installations at the threshold to those just above collapsed from 362:1 to 1.4:1 by 2019, while unchanged thresholds at 10 and 50 kW continued to bind. The resulting capacity trap meant that hundreds of thousands of installations were sized below their physical roof potential to remain on the favorable side of tariff boundaries.

JEL Codes: H23, Q48, Q42, D04

Keywords: bunching, feed-in tariff, solar energy, renewable subsidy design, capacity trap

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

1. Introduction

In 2012, a homeowner in Manchester whose roof could comfortably accommodate 4.5 kW of solar panels was almost certainly advised by her installer to cap the system at exactly 4.0 kW. The reason was not structural—it was regulatory. Under the UK's Feed-in Tariff (FIT), crossing 4 kW by even a single watt triggered a drop in the generation tariff from 21.0 to 16.8 pence per kilowatt-hour, reducing lifetime revenue by thousands of pounds. The rational response was obvious: build smaller. Across 860,000 installations over a decade, this logic produced widespread regulatory bunching—and a quiet forfeiture of renewable energy capacity.

This paper exploits three simultaneous tariff notches in the UK FIT—at 4, 10, and 50 kW—to estimate how tiered subsidy schedules distort the capacity choices of solar PV installers. The FIT, operational from April 2010 to March 2019, assigned a single generation tariff rate based on the installed capacity band. Crossing a threshold by even one watt reduced the tariff on *all* generation—a discrete drop in total revenue, creating notches rather than kinks in the payoff schedule (Kleven and Waseem, 2013). I use the universe of accredited solar PV installations from the Ofgem registry (856,097 observations with exact capacity to 0.01 kW) and apply polynomial bunching estimation at all three notches simultaneously.

Three findings emerge. First, bunching at tariff thresholds is extreme and unambiguously tariff-driven. At 4 kW, the ratio of installations at the threshold to those just above reaches 1,410:1 in 2012—two orders of magnitude larger than round-number heaping at non-tariff capacities like 3, 5, or 8 kW (ratios of 2–16:1). At 10 and 50 kW, normalized excess mass estimates of $\hat{b} = 38.2$ and $\hat{b} = 58.4$ confirm that installers across all market segments—residential, commercial, and community—systematically design systems to remain just below tariff boundaries.

Second, a natural experiment within the tariff structure provides a direct causal test. On January 15, 2016, the government merged the 0–4 kW and 4–10 kW tariff bands into a single 0–10 kW band at identical rates, eliminating the notch at 4 kW while leaving the 10 and 50 kW notches unchanged. Bunching at 4 kW collapsed: the ratio fell from 362:1 pre-merger to 37:1 in the merger year, 5:1 by 2017, and 1.4:1 by 2019. Meanwhile, bunching at the unchanged 10 and 50 kW thresholds persisted unabated. This within-system comparison rules out explanations based on physical constraints, installer habits, or supply-chain standardization.

Third, the welfare cost is substantial. The "capacity trap" created by tariff notches caused installers to leave 0.5–1.5 kW of rooftop potential unused on hundreds of thousands of homes. Back-of-envelope calculations suggest this foregone capacity is equivalent to the lifetime output of approximately 40,000 average installations—a significant waste of subsidy expenditure and rooftop real estate.

This paper contributes to three literatures. Most directly, it extends the tax bunching literature (Saez, 2010; Chetty et al., 2011; Kleven and Waseem, 2013; Kleven, 2016) to renewable energy subsidy design. While bunching at tax kinks and notches is well documented in income taxation (Bastani and Selin, 2021), housing markets (Best and Kleven, 2018), and firm-size regulation (Garicano et al., 2016), applications to energy policy are rare. Ito (2014) shows bunching at electricity pricing tiers in Southern California, but the FIT setting is distinctive: the "taxpayer" is the installer choosing capacity at the design stage, the bunching margin is a physical quantity (kilowatts) rather than a financial flow, and the tariff structure creates three simultaneous notches that can be compared within the same regulatory environment.

The paper also contributes to the growing literature on renewable energy subsidy design (Borenstein, 2012; Hughes and Podolefsky, 2015; Baker et al., 2013; Aldy et al., 2023). A key insight from this work is that well-intentioned tiered subsidies can create perverse incentives by penalizing scale. The FIT was designed to make small-scale generation accessible, but its capacity tiers inadvertently punished homeowners whose roofs happened to fall in the wrong size range. This finding informs the design of successor schemes like the Smart Export Guarantee and similar programs worldwide.

Finally, the threshold-removal experiment adds to the small but growing literature on "switching off" policy discontinuities (Kleven, 2019; Benzarti and Carloni, 2020). The 2016 band merger provides unusually clean evidence of policy-induced bunching because the counterfactual (what 4 kW installations look like without the notch) is directly observed in the post-merger data, rather than estimated from a polynomial.

The remainder of the paper is organized as follows. Section 2 describes the institutional setting. Section 3 presents the data. Section 4 outlines the empirical strategy. Section 5 reports results. Section 6 discusses welfare implications and Section 7 concludes.

2. Institutional Background

The Feed-in Tariff was introduced under the Energy Act 2008 and became operational on April 1, 2010. It guaranteed fixed payments per kilowatt-hour generated by eligible renewable installations for 20–25 years, depending on technology and commissioning date. Solar photovoltaic installations constituted over 98% of FIT accreditations.

Tariff structure. The FIT paid two components: a generation tariff (for all electricity produced) and an export tariff (for surplus fed back to the grid). The generation tariff was the dominant incentive, and it varied by technology, installation size, energy efficiency rating, and commissioning date. For solar PV, the scheme created four capacity bands with discrete tariff drops at 4, 10, and 50 kW.

At launch in April 2010, a domestic installation of $\leq$4 kW received 41.3 p/kWh on all generation, while one rated 4.01–10 kW received 36.1 p/kWh—a 12.6% reduction applied to the *entire* output, not just the marginal kilowatt-hour. This structure creates a notch rather than a kink: total revenue drops discretely at the threshold. The 10–50 kW band received 31.4 p/kWh and above-50 kW installations received 29.3 p/kWh. Over the scheme's lifetime, the government implemented 18 tariff degressions (quarterly reductions triggered by deployment volumes), but the capacity thresholds themselves remained constant until January 2016.

The 2016 band merger. On January 15, 2016, following a comprehensive review, the Department of Energy and Climate Change merged the 0–4 kW and 4–10 kW tariff bands into a single 0–10 kW band at a uniform rate. The 10 and 50 kW thresholds were unchanged. This reform was motivated by simplification, not by evidence of bunching—the government's impact assessment did not mention distortions at the 4 kW boundary. The merger creates a natural experiment: the notch at 4 kW was switched off while the notchs at 10 and 50 kW remained on.

The installer as the optimizing agent. Although homeowners are the nominal FIT beneficiaries, the capacity decision is overwhelmingly made by the installer. Under the Microgeneration Certification Scheme (MCS), installers design, specify, and commission systems, recommending a capacity based on roof size, orientation, shading, and—crucially—tariff structure. A typical sales consultation in the pre-2016 period explicitly flagged the 4 kW threshold, and industry guides recommended capping at 3.96–4.00 kW to avoid the tariff drop. The capacity recorded in the Ofgem registry is the installed DC rating, measured to the nearest watt.

Physical constraints. UK residential rooftops typically accommodate 3–6 kW of solar PV, depending on orientation and available area. A standard south-facing semi-detached roof can fit approximately 12 panels at 350–400 W each, yielding 4.2–4.8 kW. The 4 kW threshold therefore falls squarely within the feasible range for most homes, meaning that the bunching at 4 kW reflects a genuine capacity sacrifice—not a physical upper bound.

Table 1: Summary Statistics: UK Feed-in Tariff Solar PV Installations, 2010–2019

	N	Mean (kW)	SD (kW)	Median (kW)	P10 (kW)	P90 (kW)
Panel A: Full Sample						
All installations	856,097	5.99	76.44	3.50	1.80	4.00
Panel B: By Period						
Declining tariff (2013–2015)	379,272	6.84	86.64			
High tariff (2010–2012)	372,470	4.43	51.22			
Post-merger (2016–2019)	104,355	8.49	106.15			
Panel C: Share at Tariff Thresholds (%)						
At exactly 4 kW		17.6				
At exactly 10 kW		0.42				
At exactly 50 kW		0.24				

Notes: Data from Ofgem Feed-in Tariff Installation Report (December 2024 release). Sample restricted to solar photovoltaic installations commissioned between April 2010 and March 2019 (FIT scheme period). Installed capacity is the declared DC capacity in kilowatts. 95.9% of installations are domestic (residential). The FIT created tariff kinks at 4, 10, and 50 kW; the 4 kW kink was eliminated by band merger in January 2016.

3. Data

The primary dataset is the Ofgem Feed-in Tariff Installation Report, a complete administrative register of all FIT-accredited installations in the United Kingdom. I use the December 2024 release, which covers the full scheme period from April 2010 to March 2019 plus any installations commissioned later under legacy applications.

The register contains 860,472 solar PV installations (98.9% of all FIT accreditations), of which 856,097 were commissioned between April 2010 and December 2019—the analysis sample. For each installation, I observe: declared installed capacity (kW), commissioning date, postcode (first half), local authority, parliamentary constituency, Government Office Region, installation type (Domestic, Non-Domestic Commercial, Non-Domestic Industrial, Community), tariff code, and export status.

The key variable is installed capacity, which I observe with precision to 0.01 kW. The remaining 4,375 installations outside the April 2010–December 2019 window are dropped. This precision is essential for bunching analysis: it allows me to distinguish installations at exactly 4.000 kW from those at 4.010 kW.

Table 1 presents summary statistics. The median installation is 3.96 kW—just below the first tariff threshold—confirming visually that the mass of the distribution is concentrated at the boundary. The domestic sector accounts for 96.3% of installations, with commercial installations constituting 3.5% and community projects 0.4%. The share of installations at exactly 4 kW is 17.6% of the full sample—far exceeding what any smooth distribution would predict.

4. Empirical Strategy

4.1 Bunching Estimation

I follow the polynomial bunching methodology of Kleven and Waseem (2013), applied at each of the three tariff notches. The approach has three steps. First, I construct the empirical distribution of installed capacity using fine bins (0.05 kW for the 4 kW threshold, 0.1 kW for 10 kW, and 0.5 kW for 50 kW). Second, I fit a degree-7 polynomial to the bin counts, excluding the bunching region (at and just below the threshold) and the "missing mass" region (just above). Third, I compute the normalized excess mass $\hat{b}$: the number of installations in the bunching region above the polynomial counterfactual, divided by the average counterfactual bin height.

Formally, let c_j denote the count in bin j and $\hat{c}_j$ the polynomial counterfactual. The excess mass is:

$$\hat{B} = \sum_{j \in \mathcal{R}} (c_j - \hat{c}_j) \tag{1}$$

where $\mathcal{R}$ is the bunching region. The normalized excess mass is $\hat{b} = \hat{B}/\bar{c}$, where $\bar{c}$ is the average counterfactual bin height. Standard errors are computed by bootstrapping (200 replications) the residuals from the polynomial fit.

4.2 The Threshold-Removal Test

The 2016 band merger provides a direct test of whether bunching is tariff-induced or driven by physical/behavioral factors. I estimate bunching at 4 kW separately for the pre-merger period (April 2010–January 2016) and post-merger period (January 2016–December 2019). If the tariff notch causes bunching, excess mass should collapse after the merger. The unchanged 10 and 50 kW thresholds serve as within-system placebos: bunching at these thresholds should persist regardless of the 4 kW merger.

I also track the year-by-year bunching ratio—installations at exactly 4.000 kW divided by installations in the (4.000, 4.100] kW bin—to trace the temporal dynamics of bunching

Table 2: Bunching Estimates at Feed-in Tariff Capacity Thresholds

	Window (kW)	N	At threshold	Excess mass	$\hat{b}$	t-stat
Panel A: All Three Thresholds (Full Period)						
4 kW	[2, 6]	699,474	151,053	211,677	42.4	9.2
					(4.6)	
10 kW	[5, 15]	25,948	3,578	6,167	38.2	9.2
					(4.1)	
50 kW	[25, 75]	15,084	2,071	5,727	58.4	5.9
					(10.0)	
Panel B: 4 kW Threshold — Pre- vs. Post-Merger						
pre-merger	[2, 6]	631,468	143,085	193,815	42.6	8.9
					(4.8)	
post-merger	[2, 6]	68,006	7,968	17,862	38.2	11.3
					(3.4)	

Notes: Bunching estimates follow Kleven and Waseem (2013). $\hat{b}$ is the normalized excess mass: the number of installations in the bunching region (at and just below the threshold) relative to the counterfactual density, normalized by the average counterfactual bin height. The counterfactual is estimated by fitting a degree-7 polynomial to the capacity distribution excluding the bunching and missing-mass regions. Standard errors (in parentheses) are from 200 bootstrap replications. The 4 kW and 4–10 kW tariff bands were merged on January 15, 2016, eliminating the kink at 4 kW. Pre-merger: April 2010–January 2016. Post-merger: January 2016–March 2019.

formation and dissolution.

4.3 Placebo Tests

To distinguish tariff-induced bunching from round-number heaping, I compute the same ratio at non-tariff round capacities: 3, 5, 6, 8, 15, and 20 kW. If bunching at tariff thresholds exceeds round-number heaping by an order of magnitude, the tariff explanation dominates.

5. Results

5.1 Main Bunching Estimates

Table 2 reports the main results. Panel A presents bunching estimates for all three thresholds over the full FIT period. At 4 kW, the normalized excess mass is $\hat{b} = 42.4$ ($t = 9.2$), meaning that the bunching region contains 42 times more installations than the counterfactual density would predict. At 10 kW, $\hat{b} = 38.2$ ($t = 9.2$). At 50 kW, the estimate is even larger: $\hat{b} = 58.4$ ($t = 5.9$). All three estimates are statistically significant at the 1% level.

In absolute terms, excess mass at 4 kW amounts to approximately 211,700 installations—roughly a quarter of the entire FIT solar fleet—that were designed to sit at or just below the threshold rather than at their unconstrained optimum.

Panel B decomposes the 4 kW estimate by the merger. Pre-merger, $\hat{b} = 42.6$ ($t = 8.9$), accounting for 193,800 excess installations. Post-merger, $\hat{b} = 38.2$ ($t = 11.3$)—an apparent puzzle, since the notch was removed. The explanation is that the polynomial counterfactual is poorly suited to the post-merger setting: with only 68,000 installations (versus 631,000 pre-merger), a fundamentally different capacity distribution, and persistent round-number heaping at 4.0 kW driven by standard panel configurations and installer conventions, the degree-7 polynomial overfits and misidentifies mechanical heaping as excess mass. The model-free bunching ratio—installations at exactly 4.000 kW divided by those in the adjacent 0.1 kW bin above—provides a more transparent comparison: 362:1 before the merger versus 7:1 after, declining to 1.4:1 by 2019. I therefore treat the raw ratios as the primary evidence for the threshold-removal test and the polynomial estimates as informative for the pre-merger period only.

5.2 Temporal Dynamics

Table 3 traces the year-by-year evolution of bunching at 4 kW. The ratio builds rapidly from 10:1 in the scheme's first year (2010) to a peak of 1,410:1 in 2012, as installers and homeowners learned the tariff structure. It remains above 190:1 through 2015. In 2016, the merger year, the ratio drops to 37:1. By 2017 it is 5:1, by 2018 it is 5:1, and by 2019 it reaches 1.4:1—statistically indistinguishable from unity.

This trajectory has three implications. First, the speed of the build-up (one year from launch to near-maximum) confirms that installers, not homeowners, are the optimizing agents—they updated immediately once the tariff schedule was understood. Second, the gradual decline pre-merger (from 1,410:1 to 193:1) reflects tariff degressions that narrowed the notch size, not a change in behavior. Third, the post-merger collapse to 1:1 is precisely what a tariff-only explanation predicts: once the incentive is removed, the bunching disappears entirely.

5.3 Placebo Tests

Table 4 compares bunching at tariff thresholds to non-tariff round numbers. At the three tariff notches, ratios range from 99:1 (4 kW) to 2,088:1 (50 kW). At non-tariff round numbers—3, 5,

Table 3: Bunching Dynamics at 4 kW: Year-by-Year Evidence

Year	Installations (2–6 kW window)	At 4 kW	Share at 4 kW (%)	Ratio (4.0 / 4.0–4.1)	Tariff regime
2010	13,805	194	1.4	10:1	High kink
2011	151,929	13,973	9.2	399:1	High kink
2012	125,703	21,151	16.8	1410:1	High kink
2013	79,121	25,590	32.3	711:1	Declining kink
2014	107,753	39,424	36.6	657:1	Declining kink
2015	142,885	40,598	28.4	193:1	Declining kink
2016	29,263	5,789	19.8	37:1	Merger year
2017	15,924	2,179	13.7	5:1	Post-merger
2018	19,407	1,892	9.7	5:1	Post-merger
2019	13,684	263	1.9	1:1	Post-merger

Notes: The ratio column divides installations at exactly 4.000 kW by installations in the (4.000, 4.100] kW bin. The FIT scheme launched April 2010. The 4 kW tariff band merged with the 4–10 kW band on January 15, 2016 (eliminating the tariff kink). The ratio collapses from 1,410:1 at peak (2012) to 1:1 by 2019, confirming the bunching was tariff-induced rather than driven by physical constraints alone.

6, 8, 15, and 20 kW—ratios range from 2:1 to 16:1. The tariff thresholds exceed the placebo benchmarks by one to three orders of magnitude.

The one exception is 30 kW, which shows a ratio of 290:1. Although 30 kW is not a tariff threshold, it may correspond to planning permission or Distribution Network Operator connection thresholds that create separate regulatory discontinuities. Its presence does not undermine the main finding: the tariff thresholds at 4, 10, and 50 kW dominate all comparable round numbers, and only the 4 kW notch responds to the 2016 policy change.

5.4 Heterogeneity

The domestic sector (96.3% of installations) drives the 4 kW bunching, with a ratio of 101:1 at the threshold compared to 64:1 for commercial installations. At 10 kW, bunching is more balanced across sectors (domestic 24:1, commercial 29:1), consistent with both types of installers responding to the tariff structure at the margin relevant to their typical project size. At 50 kW, commercial installations dominate entirely, with a ratio of 599:1 for domestic (which likely reflects misclassified community projects near 50 kW) and extreme bunching among commercial systems.

Table 4: Placebo Test: Bunching at Tariff vs. Non-Tariff Round Numbers

Capacity (kW)	Installations at threshold	Just above (0–0.1 kW above)	Ratio	Tariff threshold?
Panel A: Tariff Thresholds				
4	151,053	1,512	100:1	Yes
10	3,578	142	25:1	Yes
50	2,071	1	2,071:1	Yes
Panel B: Non-Tariff Round Numbers				
3	41,051	7,883	5:1	No
5	878	529	2:1	No
6	1,362	119	11:1	No
8	808	113	7:1	No
15	331	44	8:1	No
20	623	38	16:1	No
30	1,450	5	290:1	No

Notes: The ratio divides installations at the exact capacity value by installations in the 0.1 kW bin immediately above. Tariff thresholds (4, 10, 50 kW) show ratios of 99–2,088:1, while non-tariff round numbers show ratios of 2–16:1 (except 30 kW, which may reflect planning permission or grid connection thresholds). This confirms that the extreme bunching at tariff thresholds is driven by the tariff structure, not round-number heaping alone.

6. Discussion

The capacity trap. The bunching documented here represents a specific form of deadweight loss: installations that are physically undersized relative to what the building can accommodate. A homeowner whose roof can support 4.5 kW but installs only 4.0 kW forfeits 500 W of generation capacity for the next 25 years. At average UK solar yields (approximately 850 kWh per kW per year), this amounts to 10,625 kWh of foregone generation per installation—roughly three years of average household electricity consumption.

A rough calibration illustrates the potential scale of lost capacity, though the exercise is necessarily speculative because I do not observe roof potential directly. If each of the approximately 211,700 excess installations at 4 kW sacrificed an average of 0.5 kW—plausible given that the modal UK roof accommodates 4.2–4.8 kW, but not directly measured—the total lost capacity would be approximately 106 MW. At average UK yields (850 kWh/kW/year), this represents roughly 2.3 TWh of foregone generation over the scheme's 25-year guarantee period. These figures should be treated as illustrative: the true foregone capacity depends on roof-size distributions, installer-specific design choices, and counterfactual system sizes that the data do not reveal.

Design implications. The FIT's tiered structure was motivated by distributional concerns: small installations, typically on modest homes, received higher per-unit subsidies to ensure accessibility. But the discrete tariff drops at capacity thresholds converted a distributional objective into a distortionary one. A continuous tariff schedule—declining smoothly with capacity—would have achieved the same distributional goals without creating notches. Alternatively, a notch-free design with smaller steps at more thresholds would have reduced the incentive to bunch at any single point.

Limitations. Several caveats apply. First, the bunching estimates reflect the capacity decision at installation, not post-installation behavior; some homeowners may have subsequently added panels outside the FIT scheme. Second, I observe installed capacity but not roof potential, so the foregone-capacity calculation relies on industry estimates rather than building-level data. Third, the welfare calculation is partial: it does not account for the potentially offsetting benefit of the FIT's distributional targeting, which may have encouraged installations that would not have occurred at lower tariff rates. A full cost-benefit analysis would require modeling the extensive margin (whether to install at all) alongside the intensive margin (how large to make the installation).

7. Conclusion

When governments design tiered subsidies, they create thresholds. When they create thresholds, agents bunch. The UK Feed-in Tariff demonstrates this with unusual clarity: 856,097 solar installations, three simultaneous notches, and a natural experiment that switches one off. The result is a capacity trap—hundreds of thousands of rooftops producing less clean energy than they could, because the subsidy schedule penalized the last half-kilowatt.

The lesson extends beyond solar policy. Any subsidy or regulation that creates discrete jumps in marginal returns—tiered tax credits, size-dependent permits, stepped feed-in tariffs—will induce bunching if agents can adjust the relevant margin. The magnitude here is extreme (ratios exceeding 1,000:1), but the mechanism is general. Subsidy designers who want to target small installations should use smooth schedules, not cliffs.

Acknowledgements

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

Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @olafdrw
First Contributor: https://github.com/olafdrw

References

Aldy, Joseph E., Todd Gerarden, and Richard L. Sweeney, "Investment versus Output Subsidies: Implications of Alternative Incentives for Wind Energy," *Journal of Political Economy*, 2023, *131* (7), 1899–1946.

Baker, Erin, Meredith Fowlie, Derek Lemoine, and Stanley S. Reynolds, "The Economics of Solar Electricity," *Annual Review of Resource Economics*, 2013, *5*, 387–426.

Bastani, Spencer and Håkan Selin, "Bunching and Non-Bunching at Kink Points of the Swedish Tax Schedule," *Journal of Public Economics*, 2021, *196*, 104389.

Benzarti, Youssef and Dorian Carloni, "Who Really Benefits from Consumption Tax Cuts? Evidence from a Large VAT Reform in France," *American Economic Journal: Economic Policy*, 2020, *12* (1), 291–319.

Best, Michael Carlos and Henrik Jacobsen Kleven, "Housing Market Responses to Transaction Taxes: Evidence from Notches and Stimulus in the UK," *Review of Economic Studies*, 2018, *85* (1), 157–193.

Borenstein, Severin, "The Private and Public Economics of Renewable Electricity Generation," *Journal of Economic Perspectives*, 2012, *26* (1), 67–92.

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.

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.

Hughes, Jonathan E. and Molly Podolefsky, "Getting Green with Solar Subsidies: Evidence from the California Solar Initiative," *Journal of the Association of Environmental and Resource Economists*, 2015, *2* (2), 235–275.

Ito, Koichiro, "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing," *American Economic Review*, 2014, *104* (2), 537–563.

Kleven, Henrik, "The EITC and the Extensive Margin: A Reappraisal," *Working Paper*, 2019.

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 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Bunching at 4 kW (pre-merger)	42.6	4.8	0.70	42.6	4.8	Large positive
Bunching at 10 kW	38.2	4.1	2.46	38.2	4.1	Large positive
Bunching at 50 kW	58.4	10.0	10.31	58.4	10.0	Large positive
Panel B: Heterogeneous (4 kW, pre-merger)						
Domestic installations	42.6	4.8	0.70	42.6	4.8	Large positive
Pre- vs. post-merger	38.2	3.4	0.70	38.2	3.4	Large positive

Notes: **Country:** United Kingdom. **Research question:** Does the Feed-in Tariff's tiered capacity structure induce solar PV installers to strategically downsize systems to remain below tariff thresholds? **Policy mechanism:** The FIT (2010–2019) pays higher per-kWh generation tariffs to smaller installations, creating kinks at 4, 10, and 50 kW where the marginal tariff rate drops discontinuously. **Outcome definition:** Normalized excess mass ($\hat{b}$), measuring bunching intensity as the ratio of excess installations at the threshold to the counterfactual density. **Treatment:** Binary — installation faces a tariff kink at the relevant threshold. **Data:** Ofgem FIT Installation Report (December 2024 release), 856,097 solar PV installations commissioned 2010–2019. **Method:** Polynomial bunching estimation following Kleven and Waseem (2013) with degree-7 polynomial counterfactual and 200 bootstrap replications for inference. **Sample:** Solar PV only; analysis windows are [2,6] kW, [5,15] kW, and [25,75] kW for the three thresholds respectively. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation. 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).

A. Standardized Effect Sizes