Kinks Without Bunching: Purchase-Tax Rate Jumps and Manufacturer CO_2 Manipulation

APEP Autonomous Research* @olafdrw

April 3, 2026

Abstract

Governments use graduated vehicle taxes to discourage CO_2 emissions, but whether manufacturers respond by manipulating type-approval emissions is unknown. I exploit four kinks in the Dutch BPM purchase tax where marginal rates jump by 1.6–39.5 times, using 1.25 million Dutch registrations and 11.5 million German registrations as a placebo. McCrary density tests find no significant discontinuity at any kink. Polynomial bunching estimates are unstable across specifications. The apparent concentration of vehicles near 79 g/km—entirely plug-in hybrids—is present equally in Germany, attributable to EU fleet-average regulation rather than national taxation. Purchase-tax kinks, unlike level notches, fail to distort manufacturer behavior.

JEL Codes: H23, L62, Q58
Keywords: vehicle taxation, bunching, CO_2 emissions, kink, purchase tax, manufacturer behavior

---

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 25m).

1. Introduction

A Toyota Corolla registered in the Netherlands at 142 g/km of CO_2 faces a one-time purchase tax of €9,767. The same car at 141 g/km pays €9,483—a €284 savings from a single gram. Across four kink points in the Dutch BPM (*Belasting van Personenauto's en Motorrijwielen*), manufacturers could save buyers hundreds of euros per gram by calibrating type-approval emissions to land just below each threshold. Whether they do so is the question this paper answers. I find no detectable evidence that they do.

A large and influential literature uses bunching at tax thresholds to estimate behavioral elasticities in settings from income taxation (Saez, 2010; Chetty et al., 2011) to housing markets (Best and Kleven, 2018). In vehicle markets, Sallee and Slemrod (2012) documented significant bunching at the US gas guzzler tax—a level notch that imposes a discrete tax increase above the threshold. Ito and Sallee (2018) showed that Japanese fuel economy standards generate bunching at weight class boundaries. These findings underpin a growing policy consensus that vehicle CO_2 regulation creates strategic gaming incentives (Reynaert, 2021; Grigolon et al., 2018). But the evidence base conflates two distinct tax instruments: notches, where the tax level jumps discontinuously, and kinks, where only the marginal rate changes. This distinction is theoretically crucial (Kleven, 2016): notches create a dominated region where no optimizing agent should locate, generating sharp bunching; kinks create only a change in slope, producing diffuse bunching that scales with the elasticity of adjustment.

This paper provides the first multi-cutoff test of manufacturer response to purchase-tax kinks in the CO_2 distribution. The Dutch BPM offers an unusually clean setting: four kinks at 79, 101, 141, and 157 g/km where the marginal tax rate increases by factors of 1.6 to 39.5. If manufacturers manipulate type-approval emissions to avoid these kinks, we should observe excess mass just below each threshold and missing mass above. The multi-cutoff structure enables a dose-response test: larger rate jumps should generate more bunching.

I use vehicle-level registration data from the European Environment Agency's CO_2 monitoring program, which records the WLTP CO_2 emissions, fuel type, manufacturer, and mass for every new car registered in the European Union. The analysis covers 1.25 million Dutch and 11.5 million German registrations from 2020–2022, the WLTP era of BPM taxation. Germany serves as a placebo country: it has no purchase tax with CO_2 kinks at the Dutch thresholds, allowing a difference-in-bunching design (Kleven and Waseem, 2013).

The results are consistently null for the Netherlands. McCrary density discontinuity tests (McCrary, 2008) find no statistically significant break at any of the four BPM kinks in the Dutch distribution ($|t| < 1$ at all four thresholds). Polynomial bunching estimates in the tradition of Chetty et al. (2011) are completely unstable: the sign of the estimated excess mass reverses across polynomial degrees and bandwidth choices at the 141 and 157 g/km kinks, and the 79 and 101 g/km kinks produce degenerate estimates due to the sharp transition between plug-in hybrid (PHEV) and conventional powertrains. Year-by-year estimates show no temporal stability.

The most striking data pattern—a concentration of vehicles near 79 g/km—turns out to be uninformative about BPM taxation. This cluster is composed entirely of PHEVs, and an identical pattern appears in German data, where no equivalent purchase-tax kink exists. The 79 g/km cluster reflects EU fleet-average CO_2 regulation under Regulation (EU) 2019/631 (European Parliament and Council, 2019), which creates strong manufacturer incentives to calibrate PHEV type-approval emissions regardless of national tax policy. In the conventional (ICE) vehicle range above 86 g/km—where BPM kinks at 101, 141, and 157 g/km should bite—the CO_2 distribution is lumpy but shows no systematic response to tax thresholds.

This null result carries specific policy implications. Graduated purchase taxes with kink-based rate structures—used in the Netherlands, Norway, Denmark, and other EU member states—are less effective at distorting manufacturer behavior than the notch-based designs studied by Sallee and Slemrod (2012) and Best and Kleven (2018). The per-gram tax savings at a kink—€79 to €568—are likely modest relative to the engineering cost of recalibrating an ICE powertrain to achieve a specific type-approval CO_2 target (Reynaert, 2021). If policymakers want to use vehicle taxation to shift the CO_2 distribution, the evidence suggests they need notches—not kinks.

The paper contributes to three literatures. First, it advances the bunching literature (Kleven, 2016) by documenting a setting where theory predicts diffuse bunching and the data confirm its absence—a rare multi-cutoff null in a high-stakes policy domain. Second, it informs the vehicle taxation literature (Anderson and Sallee, 2012; Rivers and Schaufele, 2017; Allcott and Wozny, 2014) by demonstrating that the instrument's design (kink vs. notch) dominates its level in determining manufacturer response. Third, it contributes to the debate over CO_2 manipulation in vehicle markets (Reynaert, 2021) by showing that PHEV CO_2 clustering is an EU-wide fleet regulation phenomenon, not a response to national fiscal incentives.

2. Institutional Background

The Dutch BPM. The BPM is a one-time vehicle purchase tax levied at the point of registration in the Netherlands. Since July 2020, the tax has been calculated from WLTP CO_2 emissions using a piecewise-linear schedule with five bands. Table 1 presents the full schedule. The tax is *continuous*—there is no level jump at any threshold—but the *marginal*

Table 1: Dutch BPM Tax Schedule (2020–2022, WLTP)

Band	CO_2 Range (g/km)	Base Tax	Marginal Rate	Rate Jump	Tax at Upper Bound
1	0–79	€667	€2/g	—	€825
2	80–101	€825	€79/g	39.5×	€2,563
3	102–141	€2,563	€173/g	2.2×	€9,483
4	142–157	€9,483	€284/g	1.6×	€14,027
5	158+	€14,027	€568/g	2.0×	€38,451

Notes: BPM (*Belasting van Personenauto's en Motorrijwielen*) is a one-time vehicle purchase tax in the Netherlands. The tax is a continuous, piecewise-linear function of WLTP CO_2 emissions with kinks at 79, 101, 141, and 157 g/km where the marginal rate increases. Rate Jump shows the ratio of the new marginal rate to the previous rate at each kink. Source: Wet op de belasting van personenauto's en motorrijwielen 1992, as amended 2020.

rate increases sharply at each kink. At the first kink (79 g/km), the per-gram rate rises from €2 to €79, a 39.5-fold increase. At subsequent kinks, the rate increases by factors of 2.2, 1.6, and 2.0. A vehicle at 158 g/km pays €14,595 in BPM; the same vehicle at 79 g/km pays €825.

The piecewise-linear design creates four kink points with varying incentive strengths. Unlike the US gas guzzler tax studied by Sallee and Slemrod (2012)—which imposes a discrete $1,000–$7,700 penalty above a fuel economy threshold—the BPM never creates a dominated region. A manufacturer moving a vehicle from 80 to 79 g/km saves the buyer €79 (one gram at the higher rate); moving from 158 to 157 g/km saves €568. These savings are meaningful but modest relative to vehicle prices of €25,000–€50,000.

EU fleet-average regulation. Alongside national taxation, all European manufacturers face fleet-average CO_2 targets under Regulation (EU) 2019/631. The 2021–2024 target is 95 g/km, with manufacturer-specific adjustments based on fleet mass. Exceeding the target triggers fines of €95 per g/km per vehicle sold. PHEVs, which typically achieve 30–80 g/km under WLTP, serve as compliance vehicles that pull down manufacturers' fleet averages (Reynaert, 2021). This creates a strong incentive to calibrate PHEV type-approval emissions as low as possible—independent of any national purchase tax.

Germany as placebo. Germany levies an annual CO_2-based vehicle tax (*Kfz-Steuer*) at a flat rate of €2/g above 95 g/km, with no kinks at the Dutch BPM thresholds. Any bunching at 79, 101, 141, or 157 g/km in the German distribution would indicate an EU-wide phenomenon unrelated to BPM taxation, providing a clean placebo test.

Table 2: Summary Statistics: New Vehicle Registrations (2020–2022)

	Netherlands	Germany
Total registrations	1,253,380	11,478,416
Mean CO_2 (g/km)	122.1	135.6
Median CO_2 (g/km)	124	138
Std. Dev. CO_2	34.0	46.0
10th percentile CO_2	87	46
90th percentile CO_2	157	189
Netherlands by year		
2020	518,428	
2021	502,470	
2022	232,482	

Notes: Vehicle-level registration data from the European Environment Agency, CO_2 emissions monitoring under Regulation (EU) 2019/631. CO_2 measured under the WLTP test cycle. Registrations with $CO_2 = 0$ (battery-electric vehicles) excluded. Germany serves as a placebo country: it has no purchase tax with CO_2 kinks at the Dutch BPM thresholds.

3. Data

I use vehicle-level registration data from the European Environment Agency's CO_2 monitoring program under Regulation (EU) 2019/631. Manufacturers report every new passenger car registration to the EEA with WLTP CO_2 emissions (g/km), fuel type, manufacturer, mass, engine capacity, and power. The data are publicly available via the EEA's SQL API.

The analysis sample covers new passenger car registrations in the Netherlands and Germany from 2020 to 2022—the period during which the WLTP-based BPM schedule has been in effect. I exclude battery-electric vehicles (0 g/km CO_2), which are outside all BPM bands. The final sample comprises 1,253,380 Dutch and 11,478,416 German registrations.

I construct CO_2 frequency distributions by counting registrations at each integer CO_2 value from 1 to 250 g/km, pooled across years. This is the standard data structure for bunching estimation (Kleven, 2016). I also decompose the distribution by fuel type (petrol, diesel, petrol/electric, diesel/electric) near each kink to identify powertrain composition effects.

Table 2 reports summary statistics. Dutch registrations have lower mean CO_2 (122.1 vs. 135.6 g/km in Germany), reflecting higher PHEV penetration driven by both BPM incentives and other Dutch EV subsidies. The Dutch distribution has substantial mass at both the PHEV range (30–85 g/km) and the ICE range (120–160 g/km).

4. Empirical Strategy

4.1 Bunching Estimation

I employ two complementary approaches to test for bunching at each BPM kink. The first is the standard polynomial bunching estimator in the tradition of Chetty et al. (2011) and Kleven (2016). For each kink point k, I fit a polynomial of degree p to the CO_2 frequency distribution using bins outside an excluded region $[k - h, k + h]$, predict the counterfactual density inside the excluded region, and calculate excess mass $\hat{b}$ as the difference between observed and counterfactual counts below the kink, normalized by the counterfactual height at the kink point. Standard errors are computed via Poisson bootstrap with 500 replications.

The polynomial bunching estimator is the theoretically appropriate tool for kinks, where the prediction is diffuse excess mass below the threshold rather than a sharp density break. However, the lumpiness of the CO_2 distribution (driven by popular car models clustering at specific values) makes the polynomial fit sensitive to specification choices. I therefore complement the polynomial approach with a McCrary-style density discontinuity test (McCrary, 2008), which provides a specification-free check for any break in the density. While the McCrary test is designed primarily for notches (level jumps), a kink that generates substantial bunching should also produce a visible density break at the threshold. I estimate:

$$\log(n_{\text{co2}}) = \alpha + \beta \cdot \mathbb{I}[\text{CO}_2 > k] + \gamma_1(\text{CO}_2 - k) + \gamma_2(\text{CO}_2 - k)^2 + \gamma_3 \cdot \mathbb{I}[\text{CO}_2 > k] \cdot (\text{CO}_2 - k) + \varepsilon \quad (1)$$

where n_{co2} is the count of vehicles at each integer CO_2 value within ±15 g/km of the kink. The coefficient β captures the log-density discontinuity. Under the null of no manipulation, $\beta = 0$.

4.2 Difference-in-Bunching

The key identification concern is that the CO_2 distribution may be lumpy for reasons unrelated to taxation—popular car models cluster at specific CO_2 values, and EU-wide regulation affects the distribution in all member states. I address this through a difference-in-bunching design: I estimate bunching in both the Netherlands (where BPM kinks apply) and Germany (where no equivalent kinks exist), and test whether Dutch excess bunching exceeds German excess bunching at each threshold. The difference-in-bunching parameter is:

$$\Delta\hat{b}_k = \hat{b}_k^{\text{NL}} - \hat{b}_k^{\text{DE}} \tag{2}$$

with standard error $\text{SE}(\Delta\hat{b}_k) = \sqrt{\text{SE}(\hat{b}_k^{\text{NL}})^2 + \text{SE}(\hat{b}_k^{\text{DE}})^2}$.

4.3 What This Design Can and Cannot Identify

The design identifies manufacturer response to BPM kinks under the assumption that, absent the BPM, the Dutch CO_2 distribution would be smooth (or at least as smooth as Germany's) at the kink points. The multi-cutoff structure with dose-response provides power: under the alternative hypothesis, bunching should be largest at the 79 g/km kink (39.5× rate jump) and smallest at the 141 g/km kink (1.6×). The design primarily identifies manufacturer-side manipulation (engine calibration to specific CO_2 targets), which is the channel that creates sharp bunching at integer values. Consumer-side sorting across existing models would shift demand toward lower-emission vehicles but not generate excess mass at specific gram thresholds. The BPM is paid by buyers, so manufacturer manipulation requires the tax incentive to be transmitted through pricing or competitive pressure.

5. Results

5.1 Main Results: No Density Discontinuity

Table 3 reports McCrary density discontinuity tests at all four BPM kinks. None of the four Dutch kinks shows a statistically significant density break. The t-statistics range from -0.78 (kink at 79 g/km) to 0.27 (kink at 101 g/km). For comparison, Germany shows marginally significant negative discontinuities at 79 g/km ($t = -2.25$) and 141 g/km ($t = -2.00$), both in the direction of *more* mass above the kink—the opposite of what BPM-driven bunching would predict. The difference-in-bunching estimates (NL − DE) are small and statistically indistinguishable from zero at 141 and 157 g/km.

Statistical power. The null result is informative only if the design has sufficient power to detect economically meaningful bunching. At the 141 g/km kink, the average bin count is approximately 16,000 vehicles, with standard deviation across nearby bins of roughly 3,000. Excess mass of 5% (800 vehicles) would produce a McCrary coefficient of $\log(1.05) \approx 0.049$ against a standard error of 0.13, yielding $t = 0.38$—undetectable. Excess mass of 20% (3,200 vehicles) would produce $t = 1.5$—still below conventional thresholds. The design can reliably detect excess mass of approximately 35% or more at a single kink. This implies that if

Table 3: McCrary Density Discontinuity Tests at BPM Kinks

Kink (g/km)	Rate Jump	Netherlands Disc.	Netherlands t-stat	Germany (Placebo) Disc.	Germany (Placebo) t-stat	NL−DE Δ Disc.
79	39.5×	-0.712 (0.910)	-0.78	-1.715 (0.762)	-2.25	1.003
101	2.2×	0.278 (1.023)	0.27	1.568 (0.919)	1.71	-1.290
141	1.6×	-0.078 (0.132)	-0.59	-0.152 (0.076)	-2.00	0.075
157	2.0×	-0.141 (0.190)	-0.74	-0.126 (0.087)	-1.45	-0.014

Notes: Log-density discontinuity at each BPM kink, estimated via $\log(n_{\text{co2}}) = \alpha + \beta \cdot \mathbf{1}[\text{CO}_2 > k] + f(\text{CO}_2 - k) + \varepsilon$ with a quadratic in the running variable and a slope interaction, using bins within ±15 g/km of each kink. Standard errors in parentheses. Rate Jump shows the ratio of marginal tax rates above vs. below the kink. Germany has no purchase tax with corresponding CO_2 kinks. None of the four Dutch kinks show a statistically significant density discontinuity.

manufacturers were shifting CO_2 by even 1–2 grams for a substantial fraction of vehicles, the effect would need to be concentrated within a narrow window to be detectable. The multi-cutoff structure partially compensates: if all four kinks showed consistent positive excess mass, even individually insignificant estimates would jointly constitute evidence. The data show no such pattern.

Raw density patterns. Table 4 shows raw vehicle counts and normalized densities within ±2 g/km of each kink. At 79 g/km, the Dutch density drops sharply from 0.37 per 1,000 (at 79) to 0.08 per 1,000 (at 80)—but the German density shows a nearly identical pattern (0.20 to 0.02). At 141 g/km, the Dutch and German densities are comparably smooth across the threshold. At 157 g/km, the NL/DE density ratio is stable at approximately 0.5, showing no differential response.

5.2 Instability of Polynomial Bunching Estimates

Table 5 demonstrates that the polynomial bunching estimator produces fundamentally unstable results. At the 141 g/km kink (the cleanest ICE-range kink), the sign of the difference-in-bunching estimate reverses three times across four polynomial degrees ($p = 3, 5, 7, 9$) and twice across four bandwidth choices ($h = 3, 5, 7, 10$). This instability is the signature of noise,

Table 4: CO_2 Density Near BPM Kinks: Netherlands vs. Germany

Kink	CO_2	Netherlands		Germany		NL/DE
(g/km)	(g/km)	Count	Per 1000	Count	Per 1000	Ratio
79	77	368	0.29	2,611	0.23	1.29
	78	329	0.26	1,323	0.12	2.28
	79	470	0.37	2,291	0.20	1.88 ←
	80	97	0.08	286	0.02	3.11
	81	349	0.28	560	0.05	5.71
141	139	18,368	14.65	168,111	14.65	1.00
	140	18,966	15.13	160,153	13.95	1.08
	141	14,418	11.50	179,528	15.64	0.74 ←
	142	15,148	12.09	158,010	13.77	0.88
	143	11,973	9.55	133,352	11.62	0.82
157	155	10,207	8.14	147,640	12.86	0.63
	156	7,620	6.08	100,297	8.74	0.70
	157	6,465	5.16	89,383	7.79	0.66 ←
	158	5,886	4.70	99,876	8.70	0.54
	159	6,667	5.32	90,807	7.91	0.67

Notes: Raw vehicle counts and density (per 1,000 registrations) at each CO_2 value within ±2 g/km of the three main BPM kinks (79, 141, 157 g/km). Kink at 101 omitted for space; pattern is similar. Arrows mark the kink point. NL/DE Ratio compares normalized densities across countries. If the BPM created excess bunching in the Netherlands, we would expect the NL/DE ratio to spike just below the kink and drop just above. The ratios show no such pattern at 141 or 157; the apparent drop at 79/80 is present equally in Germany.

not signal: when a genuine density discontinuity exists, bunching estimates are robust to polynomial order (Kleven, 2016). Here, the polynomial is fitting lumps from popular car models, not systematic bunching from tax optimization.

5.3 The PHEV Composition Channel

The most visually striking feature of the Dutch CO_2 distribution is a cluster of vehicles near 79 g/km. Decomposition by fuel type reveals that this cluster is composed entirely of PHEVs. Below 86 g/km, 100% of Dutch registrations are petrol/electric or diesel/electric hybrids. The first substantial mass of conventional petrol vehicles appears at 87 g/km (6,336 vehicles) and 92 g/km (11,688 vehicles), with no presence at or below the 79 g/km kink.

Table 5: Robustness: Bunching Estimates Across Specifications (Kink at 141 g/km)

	Panel A: Polynomial Degree			Panel B: Bandwidth		
Specification	$\hat{b}_{\text{NL}}$	$\hat{b}_{\text{DE}}$	Δ (t)	$\hat{b}_{\text{NL}}$	$\hat{b}_{\text{DE}}$	Δ (t)
$p = 3$	-0.31	-0.04	-0.27 (-10.7)	$h = 3$	0.23	-0.33
0.56 (23.3)						
$p = 5$	0.87	0.38	0.50 (13.1)	$h = 5$	-1.00	-0.15
-0.85 (-25.3)						
$p = 7$	-1.00	-0.15	-0.85 (-25.3)	$h = 7$	-1.37	0.99
-2.36 (-39.0)						
$p = 9$	-0.43	-0.20	-0.22 (-4.1)	$h = 10$	-0.53	-0.91
0.38 (2.6)						
Sign changes	3 of 4 specifications			2 of 4 specifications		

Notes: Normalized bunching estimates ($\hat{b}$) at the 141 g/km kink across polynomial degrees (Panel A, bandwidth $h = 5$) and bandwidths (Panel B, polynomial degree $p = 7$). Δ is the difference-in-bunching (NL − DE) with t-statistic in parentheses. The sign of $\hat{b}_{\text{NL}}$ and Δ reverses across specifications, indicating that the bunching estimate is not robust to researcher degrees of freedom. This instability is consistent with the McCrary null (Table 3): there is no genuine density discontinuity for the polynomial to detect.

This PHEV cluster is not a BPM response. Germany, which has no BPM kinks, shows an identical pattern: PHEV mass at 70–85 g/km with a sharp transition to ICE at 87 g/km. The cluster reflects EU fleet-average regulation under Regulation (EU) 2019/631, which creates powerful incentives for manufacturers to offer PHEVs with type-approval CO_2 below the fleet target. PHEV CO_2 is highly calibratable—it depends on battery size, electric range assumptions, and utility factor weighting in the WLTP test cycle—making PHEVs the instrument of fleet compliance. This EU-wide calibration dominates any marginal incentive from national purchase taxation.

5.4 Robustness

Year-by-year estimates. Bunching estimates at the 141 g/km kink range from +1.59 (2020) to −2.75 (2022) for the Netherlands and from +0.43 to −0.73 for Germany. If manufacturers were systematically calibrating to the BPM schedule, we would expect stable or increasing bunching over time as the policy became established. Instead, the estimates fluctuate wildly, consistent with noise from model-specific CO_2 variation.

PHEV-specific bunching at 79 g/km. Restricting the sample to PHEVs only and estimating bunching within the PHEV CO_2 range (50–110 g/km), I find normalized bunching of $\hat{b} = 0.087$ (SE $= 0.146$, $t = 0.59$). Even among the vehicles most susceptible to CO_2 calibration, there is no evidence of BPM-driven bunching at the 79 g/km threshold.

6. Discussion

The absence of bunching at Dutch BPM kinks carries a specific mechanistic interpretation. The key distinction is between kinks and notches. At a notch, crossing the threshold imposes a discrete tax increase—a lump-sum cost. This creates a dominated region where purchasing is strictly worse than at the threshold, generating sharp and visible bunching (Kleven and Waseem, 2013). At a kink, the tax is continuous but its slope changes. There is no dominated region; the incentive to locate below the threshold is proportional to the marginal rate change times the distance from the kink. For small adjustments (1–2 g/km), the savings are modest—€79–568 per gram—relative to the cost of powertrain re-engineering.

Engineering estimates suggest that reducing type-approval CO_2 by 1 g/km through powertrain recalibration costs manufacturers €50–150 per vehicle for ICE engines (Reynaert, 2021). At the 141 g/km kink, the marginal tax saving is €111 per gram (the rate jump from €173 to €284). In a competitive market where tax savings are fully passed through, manufacturers would manipulate only if the tax saving exceeds the engineering cost for the *marginal* gram. At the lower kinks (79 and 101), the per-gram savings (€77 and €94) may fall below the engineering cost, rationalizing the null as an economically predicted non-response rather than an absence of incentive transmission.

The contrast with Sallee and Slemrod (2012) is instructive. The US gas guzzler tax creates level jumps of \$1,000–\$7,700 at fuel economy thresholds—these are notches, not kinks—and generates clear bunching. The Dutch BPM, despite marginal rate changes of up to 39.5×, produces no observable response because the *level* of the tax is continuous. This suggests that the functional form of vehicle taxation—not its magnitude—determines whether manufacturers manipulate type-approval emissions.

For policymakers, the implication is that graduated purchase taxes with piecewise-linear rate schedules are poor instruments for influencing manufacturer CO_2 calibration. If the policy goal is to create incentives for lower-emission vehicles, a notch-based design (with discrete tax jumps at thresholds) would be more effective, as demonstrated by the UK Vehicle Excise Duty bands, the Norwegian registration tax, and the French bonus-malus system.

7. Conclusion

The Dutch BPM's four CO_2 kinks—the largest increasing the marginal tax rate by a factor of 39.5—generate no detectable bunching in the vehicle registration distribution. The apparent clustering near 79 g/km is entirely PHEVs, driven by EU fleet regulation rather than national taxation, and present identically in Germany. In this setting, purchase-tax kinks—even with dramatic rate jumps—do not generate detectable manufacturer manipulation, consistent with the absence of a dominated region that makes notches effective. Where the engineering cost of CO_2 adjustment exceeds the per-gram tax saving, the rational manufacturer response is non-response.

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

Allcott, Hunt and Nathan Wozny, "Gasoline Prices, Fuel Economy, and the Energy Paradox," *Review of Economics and Statistics*, 2014, *96* (5), 779–795.

Anderson, Soren T. and James M. Sallee, "Automobile Fuel Economy Standards: Impacts, Efficiency, and Alternatives," *Review of Environmental Economics and Policy*, 2012, *6* (1), 45–62.

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.

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.

European Parliament and Council, "Regulation (EU) 2019/631 Setting CO_2 Emission Performance Standards for New Passenger Cars and New Light Commercial Vehicles," Technical Report, Official Journal of the European Union 2019.

Grigolon, Laura, Mathias Reynaert, and Frank Verboven, "Automobile Prices, Market Structure, and New-Vehicle Fleet Composition," *Annual Review of Economics*, 2018, *10*, 295–316.

Ito, Koichiro and James M. Sallee, "Not All Regulators Are Created Equal: Fuel Economy vs. Emissions Standards," *American Economic Review*, 2018, *108* (10), 2532–2569.

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.

McCrary, Justin, "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, 2008, *142* (2), 698–714.

Reynaert, Mathias, "Abatement Strategies and the Cost of Environmental Regulation: Emission Standards on the European Car Market," *Review of Economic Studies*, 2021, *88* (1), 461–488.

Rivers, Nicholas and Brandon Schaufele, "The Effectiveness of Green Vehicle Tax Policies," *Annual Review of Resource Economics*, 2017, *9*, 53–74.

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

Sallee, James M. and Joel Slemrod, "Taxation and Market Power in the US Automobile Industry," *American Economic Review*, 2012, *102* (7), 3581–3607.

A. Data Appendix

Data source. Vehicle-level CO_2 monitoring data from the European Environment Agency, collected under Regulation (EU) 2019/631. Available at `https://discodata.eea.europa.eu`. Accessed April 2026.

Sample construction. I query registration counts at each integer CO_2 value (1–250 g/km) for each country (NL, DE) and year (2020–2022). Battery-electric vehicles ($CO_2 = 0$) are excluded. The unit of observation is a CO_2 value × country × year bin. For fuel-type decomposition, I separately query counts for petrol, diesel, petrol/electric, and diesel/electric within ±30 g/km of each kink.

BPM rates. Tax rates from the *Wet op de belasting van personenauto's en motorrijwielen* 1992, as amended by the 2020 Climate Agreement implementation. The rates in Table 1 apply to the full 2020–2022 period without change.

B. Robustness Appendix

Robustness checks are reported in the main text (Table 5). Additional details: all polynomial bunching estimates use a window of ± 25 CO_2 bins around each kink, with the excluded region defined symmetrically (h below and h above the kink). Poisson bootstrap uses 500 replications with seed 42 for replicability. McCrary tests use ±15 bins with a quadratic in the running variable and a slope interaction.

C. Standardized Effect Sizes