

Where Are All the Bunchers? Income Responses to the UK High Income Child Benefit Charge

APEP Autonomous Research* @olafdrw

March 10, 2026

Abstract

The UK's High Income Child Benefit Charge (HICBC), introduced in 2013, creates a steep taper at £50,000 of adjusted net income. Using HMRC Survey of Personal Incomes data, I find no detectable bunching in total income distributions at the threshold. Two caveats qualify this null: the SPI covers all taxpayers while HICBC affects only parents (~ 13 percent near £50,000), so moderate responses could be diluted below detection; and the SPI measures total income, not adjusted net income. Subject to these limitations, I reconcile the null with massive administrative response—740,000 families opted out by 2023—arguing the dominant margin is administrative rather than income-based. Pension contributions reduce adjusted net income without affecting total income, rendering standard bunching methods unable to detect the response. These findings suggest bunching tests may understate behavioral responses when low-cost alternatives to income adjustment exist.

JEL Codes: H24, H31, J22

Keywords: bunching, notch, child benefit, tax avoidance, adjustment margins, United Kingdom

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

In January 2013, a parent earning £55,000 per year discovered that £10,000 of her income—the stretch from £50,000 to £60,000—carried an effective marginal tax rate exceeding 60 percent, nearly double the standard higher rate. The High Income Child Benefit Charge (HICBC), introduced in the UK’s 2013/14 tax year, created one of the steepest benefit withdrawal schedules in any developed country’s tax-benefit system, clawing back 1 percent of Child Benefit for every £100 of adjusted net income (ANI) above £50,000 until the full amount is withdrawn at £60,000. For a family with two children receiving roughly £1,800 per year, this adds an implicit marginal rate of approximately 19 percentage points on top of the 40 percent income tax and 2 percent National Insurance already due in this range—pushing the combined rate above 60 percent over a £10,000 window. Economic theory predicts dramatic behavioral responses to such incentives: taxpayers should bunch just below the threshold, generating visible excess mass in the income distribution (Saez, 2010; Kleven and Waseem, 2013). Yet when I examine the income distribution of UK taxpayers using HMRC’s Survey of Personal Incomes, the bunching is not detected in total income distributions.

This paper applies bunching estimation methods (Kleven and Waseem, 2013; Kleven, 2016) to study behavioral responses to the HICBC notch. I construct income density distributions from 99 percentile points of the SPI and from 11 percentile points of the Annual Survey of Hours and Earnings (ASHE), estimate counterfactual densities using polynomial fitting outside an exclusion window around £50,000, and measure excess mass ratios year by year from 2005 to 2022. The central finding is a null result: the mean excess mass ratio in the post-HICBC period (2013–2022) is $\hat{b} = -0.023$ (SE = 0.038), statistically indistinguishable from the pre-HICBC baseline of $\hat{b} = -0.024$ (SE = 0.039), with a pre–post difference of 0.001 (difference SE = 0.054, $p = 0.995$).

This null result is striking because it coexists with overwhelming administrative evidence of behavioral response. By 2024, 712,000 families had opted out of receiving Child Benefit entirely to avoid the charge, up from 43,000 in the first year. Child Benefit take-up fell from 97 percent to 88 percent. HMRC collected £525 million in HICBC revenue in 2022/23 from 440,000 liable individuals—families who continued claiming but whose highest earner exceeded the threshold. These numbers represent a massive behavioral disruption to a universal benefit that had operated for decades.

Before proceeding, two data limitations bear directly on how this null should be interpreted. First, the SPI covers all UK income taxpayers, while the HICBC applies only to the higher earner in a household that is actively claiming Child Benefit. At any income level near £50,000, Child Benefit claimants represent only a fraction of all taxpayers—approximately 13

percent, based on roughly 1.1 million families with a highest earner near that threshold out of some 8 million total taxpayers in that income range. A behavioral response of $b = 0.2$ among treated families, if those families constitute 13 percent of the local population, would appear as $b \approx 0.026$ in the all-taxpayer distribution—potentially within our confidence intervals of ± 0.07 – 0.10 . Second, the SPI records total income before tax, not adjusted net income (ANI), the income concept to which the HICBC applies. These two gaps—treatment dilution and running variable mismatch—mean that our null finding is consistent with either no total-income response among affected families, or a moderate response that is diluted or shifted in ways the data cannot capture. Both limitations are inherent to working with published aggregate statistics rather than individual-level administrative microdata.

I argue that this divergence between the statistical bunching null and the administrative response reveals something fundamental about the nature of behavioral responses to tax notches. The bunching framework assumes that the primary margin of adjustment is income itself: taxpayers reduce reported or actual income to avoid crossing the threshold. But when cheaper adjustment margins exist—particularly administrative options like forgoing the benefit, or deductions that narrow the tax base without reducing total income—the income distribution may show no response even as behavior changes dramatically.

In the HICBC context, three channels are available to affected families. First, they can reduce their *adjusted net income* (ANI) below £50,000 through pension contributions, which reduce ANI but not total income as measured in HMRC’s Survey of Personal Incomes. Second, they can opt out of receiving Child Benefit entirely, avoiding the charge without any income adjustment. Third, they can reduce actual labor supply or earnings. The first two channels are invisible to standard bunching analysis of total income distributions, and the administrative data is consistent with administrative channels dominating.

To decompose these channels, I compare bunching patterns in two datasets. The SPI captures all taxpayers including the self-employed, while ASHE covers only PAYE employees. If responses operate through salary sacrifice (a pension contribution mechanism available primarily to employees), ASHE should capture them. If responses operate through self-employment income manipulation, they should appear in the SPI but not ASHE. I find that neither dataset shows significant bunching, but the residual (SPI minus ASHE) shifts from approximately zero pre-reform to a modestly positive value post-reform, consistent with a small non-PAYE response emerging over time. The magnitude remains economically small relative to the administrative response.

This paper contributes to three literatures. First, it extends the bunching literature (Saez, 2010; Kleven and Waseem, 2013; Chetty et al., 2011; Bastani and Selin, 2014) by documenting a case where a powerful notch generates minimal income bunching, highlighting

the importance of the adjustment margin for interpreting bunching estimates. Most bunching studies examine settings where the relevant income concept aligns with the bunching variable (e.g., taxable income kinks). The HICBC is unusual because the notch applies to adjusted net income while the observable distribution is total income—a distinction that matters when deductions provide a pathway between the two.

Second, it contributes to the literature on the elasticity of taxable income ([Feldstein, 1999](#); [Gruber and Saez, 2002](#); [Kopczuk, 2005](#); [Saez et al., 2012](#); [Adam and Miller, 2021](#)), which emphasizes that behavioral responses depend on the tax base definition. [Slemrod \(1996\)](#) distinguishes avoidance and real responses, arguing that the welfare costs of the two differ substantially; [Saez et al. \(2012\)](#) survey this literature and emphasize that observed elasticities conflate real and avoidance margins. When taxpayers can adjust the base through deductions rather than changing real economic activity, the welfare implications differ fundamentally ([Chetty, 2009, 2012](#)). The HICBC provides a clean case: the same income level triggers the charge, but the cheapest response (pension saving) is welfare-improving rather than distortionary.

Third, it informs policy design for benefit withdrawal. The Mirrlees Review ([Mirrlees et al., 2011](#)) recommended integrating tax and benefit systems to avoid extreme marginal rates. The HICBC violates this principle spectacularly, yet the behavioral response is primarily administrative rather than allocative. The 2024 reform raising the threshold to £60,000 provides a natural experiment—if bunching appears at £60,000 in future data, it would suggest that the response was building but needed higher income levels (where pension avoidance is harder) to manifest in total income.

2. Institutional Background

2.1 Child Benefit in the UK

Child Benefit is the UK’s principal universal family allowance, paid to families with children under 16 (or under 20 if in approved education). Introduced in its modern form in 1977, it replaced the Family Allowance and child tax allowance, consolidating support into a single non-means-tested payment. As of 2022/23, the benefit pays £21.80 per week for the eldest child and £14.45 for each additional child, amounting to £1,133 annually for one child or £1,885 for two. The payment is made to the person primarily responsible for the child’s care (historically, the mother) and is not taxable.

Until 2013, Child Benefit was genuinely universal: all families with qualifying children received it regardless of income. Take-up rates exceeded 97 percent, and the benefit was widely viewed as one of the most efficient elements of the welfare state, with low administrative

costs and near-complete coverage (Browne and Hood, 2012). Universality was a deliberate design choice: by paying the benefit to everyone regardless of means, the system avoided the stigma and low take-up associated with means-tested programs, and ensured that children in all households—including those where family income is volatile or where one partner controls household finances—received direct support.

The political economy of Child Benefit’s universality is important context for understanding the HICBC. During the 2010 austerity program, the Conservative-Liberal Democrat coalition sought visible savings from higher-income households. Rather than formally means-testing Child Benefit (which would have required primary legislation and generated administrative costs for the entire caseload), the government introduced the HICBC as a tax charge on the higher-earning partner. This preserved the legal universality of Child Benefit—anyone can still claim it—while clawing it back through the tax system. The design was politically expedient but created unusual incentives, as we shall see.

A distinctive feature of the HICBC is that it applies to *individual* income, not household income. A couple where both partners earn £49,000 (combined £98,000) faces no charge. A single-earner couple where one partner earns £55,000 (combined £55,000) loses half their Child Benefit. This creates horizontal inequity that has been widely criticized and generates incentives for income splitting between partners where feasible. The individual basis also means that the charge is determined by the highest-earning partner’s adjusted net income, creating a sharp individual-level notch rather than the smoother phase-outs typical of household-based benefit withdrawals (Brewer et al., 2010).

2.2 The High Income Child Benefit Charge

The HICBC was announced in October 2010 and took effect on 7 January 2013. It applies when the higher-income partner in a household has adjusted net income (ANI) exceeding £50,000. For ANI between £50,000 and £60,000, the charge is 1 percent of the full Child Benefit amount for each £100 of income above £50,000. At £60,000 and above, the charge equals the full Child Benefit, effectively withdrawing it entirely.

This creates a notch in the budget constraint. Consider a family with two children receiving £1,885 in annual Child Benefit. A parent earning exactly £50,000 faces no charge. A parent earning £50,100 owes 1 percent of £1,885 = £18.85. The implied marginal tax rate on the £100 increment is $(18.85 + 40 + 2)/100 = 60.85$ percent, combining the 40 percent higher-rate income tax, 2 percent National Insurance, and the HICBC withdrawal.¹ Over the

¹Above £50,270, the higher rate of income tax applies at 40 percent, and Class 1 employee NICs fall to 2 percent. The combined rate including HICBC withdrawal thus exceeds 60 percent through the withdrawal range.

full £10,000 withdrawal range, the total HICBC amounts to the full Child Benefit—creating a substantial implicit tax wedge.

Crucially, the HICBC applies to *adjusted net income*, not gross income. ANI is calculated by taking total income and subtracting pension contributions, Gift Aid donations, and certain other reliefs. This means that a taxpayer earning £55,000 in gross income who makes £5,000 in pension contributions has an ANI of £50,000 and faces no HICBC charge. Pension contributions thus provide a direct and legal mechanism to avoid the charge while simultaneously building retirement savings.

2.3 Administrative Response

The behavioral response to the HICBC has been dramatic, though not through the channel that standard bunching analysis would detect. Families have three options when the higher earner’s income approaches £50,000:

1. **Opt out of Child Benefit:** Families can elect not to receive the benefit, avoiding the need to file a self-assessment tax return and pay the charge. By August 2024, 712,000 families had opted out, representing approximately 8 percent of all Child Benefit claimant families.
2. **Claim and pay the charge:** Families can continue receiving Child Benefit and pay the HICBC through self-assessment. In 2022/23, 440,000 individuals were liable, generating £525 million in revenue.
3. **Reduce adjusted net income below £50,000:** Through pension contributions, Gift Aid, or other reliefs, taxpayers can bring their ANI below the threshold.

The scale of the administrative response is shown in [Figure 4](#). Opt-outs grew rapidly from 43,000 in the first year (2013) to over 700,000 by 2023, while Child Benefit take-up fell from 97 percent to 88 percent. The 2024 reform raising the threshold to £60,000 led to a slight decline in opt-outs to 712,000, suggesting some families returned to claiming.

2.4 The 2024 Reform

In the Spring Budget 2024, the Chancellor announced that the HICBC threshold would rise from £50,000 to £60,000 effective April 2024, with the withdrawal range extending to £80,000 ([Office for Budget Responsibility, 2024](#)). This reform effectively removed the notch at £50,000 and created a new, shallower withdrawal between £60,000 and £80,000. The reform will provide a natural falsification test when post-2024 income data become

available: if behavioral bunching was occurring at £50,000 in the income distribution, it should dissipate after April 2024 as the notch moves. The analysis in this paper uses SPI data through 2022/23, predating the reform.

3. Conceptual Framework

3.1 Bunching at Notches

Following [Kleven and Waseem \(2013\)](#), consider a taxpayer with ability n choosing income z to maximize utility $u(c, z; n)$ where $c = z - T(z)$ is consumption and $T(z)$ is the tax-benefit schedule. A notch at income z^* creates a discontinuous jump in the tax liability: $T(z^* + \epsilon) = T(z^*) + \Delta$ for some positive charge Δ .

In the standard model, all taxpayers with counterfactual income (the income they would choose absent the notch) between z^* and some upper bound $z^* + \Delta z$ will find it optimal to bunch at exactly z^* . The excess mass at z^* relative to the counterfactual density identifies the behavioral response:

$$\hat{b} = \frac{\text{Excess mass at } z^*}{\text{Counterfactual mass at } z^*} = \frac{\int_{z^*}^{z^* + \Delta z} f_0(z) dz}{f_0(z^*) \cdot h} \quad (1)$$

where $f_0(z)$ is the counterfactual density and h is the bin width.

3.2 The HICBC as a Quasi-Notch

Standard bunching theory, following [Kleven and Waseem \(2013\)](#), characterizes a notch as a discrete jump in the budget constraint at a single income point z^* . The HICBC departs from this canonical form in an important way. Rather than withdrawing Child Benefit all at once at £50,000, it phases it out at 1 percent per £100 over the £50,000–£60,000 range—a structure more precisely described as a steep taper or quasi-notch. The effective marginal rate increase is approximately 19 percentage points for a two-child family (the full Child Benefit divided by the £10,000 withdrawal range), which is large but spread over a £10,000 window rather than concentrated at a single point.

This distinction matters for the empirical analysis. In a pure notch, bunching should be highly concentrated just below £50,000. In a quasi-notch with a £10,000 phase-out range, the incentive to bunch begins at £50,000 but a taxpayer earning £52,000 still faces a substantial marginal rate and retains an incentive to reduce income. The bunching response, if any, should therefore be spread over a broader region—the entire £45,000–£60,000 range—rather than forming a sharp spike. At the same time, a £10,000 phase-out is still far shorter than

a gradual taper typical of means-tested benefits; the implicit rate of 19pp is still nearly half the income tax rate, and the incentive to cross *below* £50,000 entirely is large. The quasi-notch thus shares the key feature of a pure notch—a discrete benefit from reaching the entry threshold—even if the budget constraint is not literally discontinuous at a single point. Our estimation window of £45,000–£55,000 captures the lower half of the phase-out range and should be wide enough to detect bunching if it exists.

3.3 Multiple Adjustment Margins

The standard framework assumes a single adjustment margin: reported income. In the HICBC setting, taxpayers have access to at least three margins:

1. **Real income adjustment** ($z \rightarrow z'$): Reducing hours, declining promotions, or shifting income between periods.
2. **Tax-base adjustment** ($ANI \rightarrow ANI'$): Increasing pension contributions or Gift Aid to reduce ANI below £50,000 without changing total income.
3. **Administrative avoidance**: Opting out of Child Benefit entirely, making the notch irrelevant.

Let c_1, c_2, c_3 denote the costs of these three margins. For most affected families, $c_3 < c_2 < c_1$: opting out is free (merely declining a payment), pension contributions have a financial cost but also provide a return, and real income reduction is most costly. Rational taxpayers will exhaust cheaper margins before resorting to more expensive ones.

Critically, only margin 1 generates bunching in the *total income* distribution. Margin 2 generates bunching in the *ANI* distribution (which is not publicly available at fine resolution), and margin 3 generates no bunching at all. Standard bunching analysis of total income therefore captures only the most expensive adjustment channel—the one that theory predicts will be used last.

3.4 Testable Predictions

This framework generates several predictions:

1. **Minimal bunching in total income**: If administrative exit and pension avoidance are sufficiently cheap, few taxpayers will adjust total income.
2. **Large administrative response**: Opt-out rates should be substantial, especially among families with modest benefit amounts relative to the hassle cost.

3. **Channel decomposition:** Any bunching that does occur should be concentrated among taxpayers with limited access to pension deductions—primarily the self-employed, who cannot use salary sacrifice schemes.
4. **Gradual emergence:** As awareness of the HICBC grows and financial advisers optimize, the pension contribution channel may strengthen over time, but this would appear in ANI bunching, not total income bunching.

4. Data

I combine four data sources to construct the analysis datasets.

4.1 HMRC Survey of Personal Incomes (SPI)

The SPI is HMRC’s primary source for the distribution of personal incomes in the UK. I use Table 3.1a, which reports 99 percentile points (P1 through P99) of total income before tax for each tax year from 1999/2000 to 2022/23 ([HM Revenue and Customs, 2024](#)). These percentile points allow construction of a piecewise-constant density estimate of the income distribution.

The SPI covers all taxpayers including employees, the self-employed, pensioners, and those with investment income. Crucially, the income concept is *total income before tax*, not adjusted net income. Pension contributions, Gift Aid, and other deductions that reduce ANI are not subtracted. This means that taxpayers who use pension contributions to bring their ANI below £50,000 will still appear at their pre-deduction income level in the SPI distribution.

From the 99 percentile points, I construct density estimates following a quantile-function differentiation approach. For adjacent percentiles p_j and p_{j+1} with corresponding incomes z_j and z_{j+1} :

$$\hat{f}(z) = \frac{(p_{j+1} - p_j)/100}{z_{j+1} - z_j} \quad (2)$$

This yields approximately 100 bins per year, with bin widths of roughly £1,500–£3,000 in the £30,000–£80,000 range relevant for the HICBC analysis. While coarser than individual-level microdata, this resolution is sufficient to detect the type of bunching that notches generate, which involves substantial mass movements over a wider range than kink-point bunching ([Kleven, 2016](#)).

I add boundary points at $P_0 = 0$ and $P_{100} = 1.5 \times P_{99}$ to close the distribution. The raw SPI data contains 2,277 percentile-year observations (99 percentile points \times 23 tax years);

the boundary points are added during density construction. One tax year is excluded from the bunching analysis: the 2008/09 year is not reported in the SPI percentile tables for this period and is therefore omitted from all year-by-year estimates. The pre-HICBC sample thus covers seven tax years (2005, 2006, 2007, 2009, 2010, 2011, 2012).

4.2 Annual Survey of Hours and Earnings (ASHE)

ASHE is the UK’s primary source for the distribution of employee earnings, based on a 1 percent sample of PAYE (Pay As You Earn) records ([Office for National Statistics, 2024](#)). I access ASHE percentile data through the NOMIS API, obtaining the 10th, 20th, 25th, 30th, 40th, 50th, 60th, 70th, 75th, 80th, and 90th percentiles of annual gross earnings for full-time employees, separately by sex, from 2002 to 2024. However, the £50,000 threshold only falls within an evaluable ASHE bin from 2009 onward—in earlier years, the gap between adjacent percentiles spanning £50,000 is too wide for meaningful density estimation. The main bunching analysis therefore uses ASHE data from 2009 through 2022.

ASHE captures only PAYE employees—it excludes the self-employed, company directors paying themselves through dividends, and those with significant non-employment income. This makes it a natural comparison group for the SPI: if bunching appears in the SPI but not ASHE, the response is driven by non-PAYE taxpayers (the self-employed and those with investment income) who face lower adjustment frictions.

The ASHE data provides 11 percentile points per year, yielding coarser density estimates than the SPI. Density construction follows the same approach as [Equation \(2\)](#).

4.3 HMRC Income Distribution Tables

I supplement the percentile data with HMRC Tables 3.5 (pension contributions by income band) and 3.11 (self-employment income by income band) from the collated personal income statistics ([HM Revenue and Customs, 2024](#)). Table 3.5 provides the number of taxpayers and pension relief amounts in income bands of approximately £10,000 width, allowing me to assess the prevalence of pension contributions near the £50,000 threshold. Table 3.11 provides analogous data for self-employment income.

4.4 Child Benefit Administrative Data

I compile administrative statistics on Child Benefit from published HMRC releases ([HM Revenue and Customs, 2023](#)). These include: families opted out of Child Benefit (2013–2024), individuals liable for HICBC (2018–2022), HICBC revenue (2018–2022), and Child Benefit take-up rates (2013–2023). These administrative data extend beyond the SPI/ASHE income

data (which end in 2022/23) and document the scale of the administrative response to the HICBC, including the post-2024 reform period.

4.5 Summary Statistics

Table 1 presents key characteristics of the income distribution and the HICBC policy environment.

Table 1: Summary Statistics

Variable	Value
<i>Panel A: Income Distribution (SPI)</i>	
Total UK income taxpayers, 2022/23 (millions)	34.5
P80 total income, 2012 (pre-HICBC)	£36,700
P80 total income, 2022 (post-HICBC)	£47,900
P85 total income, 2012	£41,800
P85 total income, 2022	£53,700
<i>Panel B: HICBC Policy Parameters</i>	
HICBC threshold, 2013/14–2023/24	£50,000
HICBC threshold, 2024/25–present	£60,000
Withdrawal rate	1% per £100
Full withdrawal at	£60,000 (2013/14–2023/24); £80,000 (2024/25–)
<i>Panel C: Administrative Response</i>	
Individuals liable for HICBC, 2022/23 (thousands)	440
HICBC revenue, 2022/23 (£millions)	525
Families opted out of CB, 2024 (thousands)	712
CB take-up rate, 2012 (%)	97
CB take-up rate, 2022 (%)	88
<i>Panel D: Pension Contributions Near Threshold</i>	
Taxpayers £40k–£50k with pension (000s)	2,250
% with pension, £40k–£50k band	61.1
Mean pension relief, £40k–£50k	£5,733
Taxpayers £50k–£70k with pension (000s)	2,050
% with pension, £50k–£70k band	62.7
Mean pension relief, £50k–£70k	£8,634

Notes: Panel A: HMRC SPI Table 3.1a. Panel B: statutory parameters. Panel C: published HMRC statistics. Panel D: HMRC Table 3.5.

Two features of Table 1 are worth highlighting. First, the 80th percentile of total income was £36,700 in 2012 and £47,900 in 2022, meaning the £50,000 HICBC threshold sits between

the 80th and 85th percentile of the income distribution—squarely in the upper-middle range. This is a densely populated region, which should provide statistical power for detecting bunching. Second, Panel D shows that over 60 percent of taxpayers in the relevant income range already have pension contributions, with mean relief of £5,700–£8,600. These taxpayers already have the infrastructure to increase contributions to bring ANI below £50,000, making the pension channel readily available.

5. Empirical Strategy

5.1 Bunching Estimation

I estimate excess mass at the £50,000 threshold following the approach of [Kleven and Waseem \(2013\)](#) and [Kleven \(2016\)](#). The procedure has three steps.

Step 1: Counterfactual density. For each tax year, I fit a polynomial of degree K to the log density in the income range £25,000–£100,000, excluding an exclusion window of \pm £5,000 around the £50,000 threshold:

$$\ln \hat{f}(z_j) = \sum_{k=0}^K \alpha_k z_j^k + \epsilon_j, \quad z_j \notin [45,000, 55,000] \quad (3)$$

where z_j is the midpoint of density bin j . The fitted values from this regression provide the counterfactual density $\hat{f}_0(z)$ within the exclusion window.

Step 2: Excess mass. I compute the excess mass ratio \hat{b} as the integrated difference between observed and counterfactual density within the exclusion window, normalized by the counterfactual mass:

$$\hat{b} = \frac{\sum_{j \in W} (\hat{f}(z_j) - \hat{f}_0(z_j)) \cdot \Delta z_j}{\sum_{j \in W} \hat{f}_0(z_j) \cdot \Delta z_j} \quad (4)$$

where $W = \{j : z_j \in [45,000, 55,000]\}$ and Δz_j is the bin width.

Step 3: Inference. Standard errors are computed using a residual bootstrap with 500 replications. In each replication, I resample the residuals from the polynomial fit, reconstruct the density, and recompute \hat{b} . The standard error is the standard deviation across bootstrap replications.

5.2 Baseline Specification

The baseline specification uses polynomial degree $K = 7$, an exclusion window of £45,000–£55,000, and an estimation range of £25,000–£100,000. I report results for individual tax

years and for pooled pre-HICBC (2005–2012) and post-HICBC (2013–2022) periods.²

5.3 Channel Decomposition

To decompose the response across adjustment margins, I compare bunching estimates from two datasets:

- **SPI-based bunching** (\hat{b}^{SPI}): Captures all taxpayers. Reflects real income adjustment by anyone plus self-employment income manipulation. Does *not* reflect pension contributions (which reduce ANI but not total income).
- **ASHE-based bunching** (\hat{b}^{ASHE}): Captures PAYE employees only. Reflects salary sacrifice (a pension mechanism) only if it reduces gross pay. Does *not* reflect self-employment manipulation.

The difference $\hat{b}^{\text{SPI}} - \hat{b}^{\text{ASHE}}$ captures the non-PAYE channel: self-employed and investment income adjustment that appears in total income.

5.4 Identification Assumptions

The bunching approach requires that the counterfactual income distribution is smooth through the threshold. This is plausible because:

1. The £50,000 threshold does not coincide with any other major policy discontinuity during most of the analysis period. The higher-rate income tax threshold was £34,370 in 2012/13 and rose gradually to £50,270 by 2022/23.³
2. Round-number bunching at £50,000 predates the HICBC, so I use a difference-in-differences approach comparing pre- and post-2013 bunching to difference out any time-invariant heaping.

²Throughout, year labels refer to the *starting* year of UK tax years: “2012” denotes 2012/13 (April 2012–April 2013) and “2013” denotes 2013/14. The HICBC took effect on 7 January 2013, making 2012/13 a partially treated year (the charge applied only to Child Benefit received from January to April 2013). I include 2012/13 in the pre-HICBC period because the charge covered only the final quarter and any behavioral response to a partial-year charge would be attenuated. Reclassifying 2012/13 as post-HICBC or excluding it does not change the results: the pre-period mean without 2012 is -0.011 ($SD = 0.12$) and the post-period mean is unchanged at -0.023 .

³The higher-rate threshold reaches the neighborhood of £50,000 only in the final years of the sample (2021–2022). However, the higher-rate threshold is a kink (a marginal rate change), not a notch, and its proximity would bias *toward* finding bunching—taxpayers facing two incentives to cluster below £50,000 should bunch more, not less. The null finding is therefore conservative with respect to this overlap: the presence of two reinforcing incentives to cluster below £50,000 should make bunching *easier* to detect, yet none appears.

3. Placebo tests at other round numbers (£40,000, £45,000, £55,000, £60,000) confirm that the bunching estimate at £50,000 is not anomalous.

5.5 Treatment Dilution and Power

A central limitation of this analysis is that the SPI covers all UK income taxpayers, while the HICBC applies only to the higher earner in a household that is actively claiming Child Benefit. This creates a treatment dilution problem: even if HICBC-affected families exhibit substantial bunching in their income distribution, the signal may be too small to detect in the all-taxpayer distribution.

To quantify this concern, consider the following back-of-envelope calculation. In 2022/23, approximately 1.1 million families had a highest earner near the £50,000 threshold (combining the 440,000 liable for the charge with a similar number who had already opted out or who fell just below the threshold). The total number of taxpayers with income between £40,000 and £60,000 in 2022/23 was approximately 8 million, based on the SPI’s reported income distribution. This implies that Child Benefit claimants represent roughly 13–14 percent of taxpayers in the relevant income range. If the true excess mass ratio among the treated population is $b_{\text{treated}} = 0.20$ (a moderate response by the standards of the literature), the all-taxpayer excess mass would be:

$$b_{\text{all}} \approx 0.13 \times b_{\text{treated}} + 0.87 \times 0 \approx 0.026$$

This is well within the confidence intervals reported in [Table 2](#) (SEs of 0.038–0.054 for pooled means). A response of $b_{\text{treated}} = 0.50$ would produce $b_{\text{all}} \approx 0.065$, still within our confidence bounds. Only a very large response ($b_{\text{treated}} \gtrsim 0.75$) would be detectable at conventional significance levels in the all-taxpayer distribution.

The running variable mismatch compounds this problem. The HICBC threshold is defined in terms of ANI, but the SPI records total income before tax. Taxpayers who use pension contributions to bring their ANI below £50,000 will appear in the SPI at their pre-deduction total income, which may be £52,000, £55,000, or higher. These taxpayers would *not* generate bunching at £50,000 in the SPI even if they are optimally responding to the policy. The pension contribution channel therefore bypasses the bunching estimator entirely.

Taken together, these two limitations—treatment dilution and running variable mismatch—mean that our null finding is most accurately characterized as: “we do not detect bunching in the all-taxpayer total income distribution at £50,000.” This is consistent with either (a) no total-income response among affected families, (b) a moderate total-income response that is diluted below detection, or (c) a substantial response operating through pension contributions

that shifts mass in the ANI distribution without moving total income. The administrative evidence discussed in [Section 6](#) supports the latter two interpretations.

A formal power analysis under scenario (b) suggests that, to detect $b_{\text{all}} = 0.05$ with 80 percent power given our estimated SE of ≈ 0.04 for annual estimates, we would need a standard error of approximately 0.025—roughly four times as many independent observations as the 10 post-reform years we have, assuming a similar cross-year variance. Detection would require either individual-level microdata (which would also resolve the running variable mismatch) or a longer post-reform panel.

5.6 Threats to Validity

The main threats are:

Data granularity. The SPI provides 99 percentile points, yielding bins of approximately £2,000–£3,000 near £50,000. This may miss very narrow bunching within a single bin. However, notch responses typically generate bunching over a wider range than kink responses ([Kleven, 2016](#)), making this concern less severe.

Income concept mismatch. The SPI reports total income, not ANI. Bunching in ANI through pension contributions is unobservable. This is a genuine limitation: the policy-relevant income concept and the measurable income concept diverge in a way that systematically prevents detection of the cheapest avoidance channel. Documenting this mismatch is part of the paper’s contribution, but it does not eliminate the measurement gap ([Saez et al., 2012](#)).

Composition effects. Nominal income growth moves the fraction of taxpayers near £50,000 over time, potentially confounding pre-post comparisons. I address this by examining year-by-year estimates rather than relying solely on pooled comparisons.

6. Results

6.1 Main Bunching Estimates

[Figure 1](#) presents the key visual evidence. The figure plots average income density in the £30,000–£80,000 range for the pre-HICBC period (2005–2012, dashed blue) and the post-HICBC period (2013–2022, solid red), with a vertical line at the £50,000 threshold. If the HICBC generated bunching in total income, the post-HICBC density should show a visible spike or bulge just below £50,000 relative to the pre-HICBC counterfactual. No such pattern is apparent. The two density curves track each other closely throughout the range, with the post-HICBC density shifted rightward reflecting real income growth.

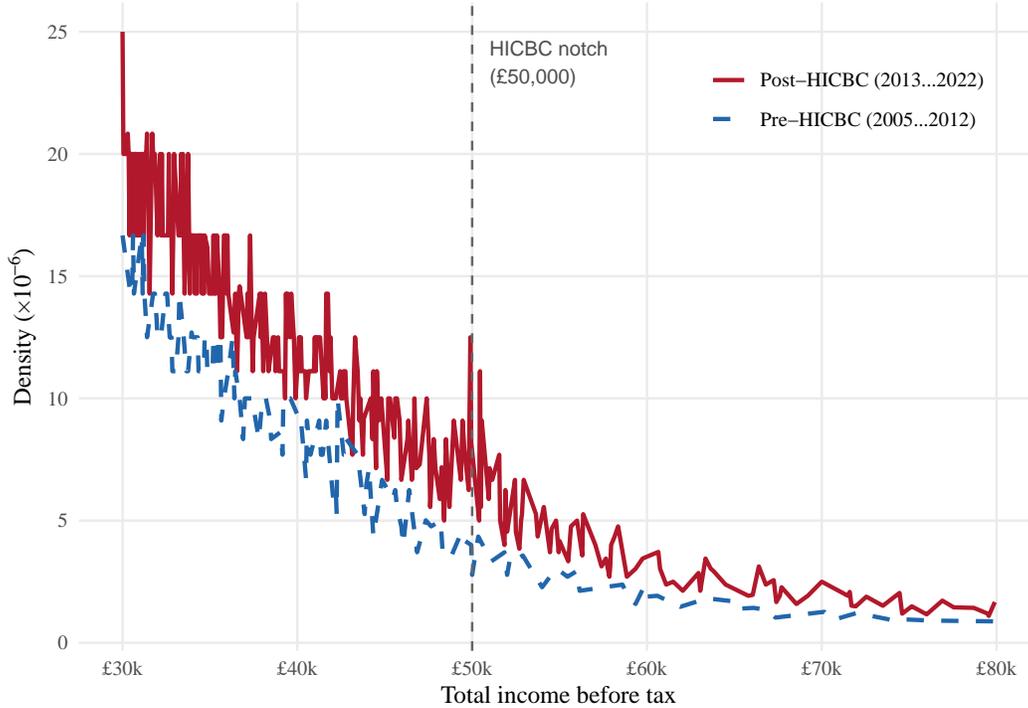


Figure 1: Income Density Near the HICBC Threshold, Pre- vs. Post-Reform
Notes: Average density of total income before tax, computed from 99 SPI percentile points. Pre-HICBC average is 2005–2012; post-HICBC is 2013–2022. Vertical dashed line marks the £50,000 HICBC threshold. Density is measured in units of 10^{-6} per pound.

The year-by-year estimates in [Table 2](#) confirm the visual evidence: in no post-HICBC year does the bunching estimate achieve statistically significant positive excess mass. The pre-HICBC mean is $\hat{b} = -0.024$ and the post-HICBC mean is $\hat{b} = -0.023$. The difference of 0.001 is economically negligible and statistically insignificant. Several individual years show statistically significant *negative* bunching (2010, 2011, 2014, 2015, and 2016, all with $p < 0.05$). Among positive post-HICBC estimates, only 2021 reaches marginal significance ($\hat{b} = 0.128$, $p = 0.029$), while 2017, 2018, 2019, 2020, and 2022 remain insignificant.

Table 2: Bunching Estimates at £50,000: Full Time Series

Tax Year	\hat{b}	SE	t -stat	p -value
<i>Pre-HICBC</i>				
2005	0.189	0.173	1.09	0.275
2006	-0.052	0.053	-0.99	0.322
2007	0.026	0.063	0.40	0.686
2009	-0.059	0.034	-1.75	0.081
2010	-0.101	0.033	-3.08	0.002
2011	-0.100	0.046	-2.18	0.030
2012 [†]	-0.068	0.040	-1.71	0.087
<i>Post-HICBC</i>				
2013	-0.070	0.059	-1.19	0.234
2014	-0.154	0.049	-3.12	0.002
2015	-0.155	0.041	-3.75	<0.001
2016	-0.189	0.047	-4.00	<0.001
2017	-0.087	0.049	-1.77	0.078
2018	0.056	0.050	1.14	0.256
2019	0.101	0.063	1.61	0.107
2020	0.073	0.044	1.66	0.096
2021	0.128	0.059	2.19	0.029
2022	0.062	0.052	1.19	0.234
<i>Pooled means^a</i>				
Pre-HICBC (2005–2012)	-0.024	[0.039]		
Post-HICBC (2013–2022)	-0.023	[0.038]		
Difference	0.001	[0.054]	0.01	0.995

Notes: Excess mass ratio \hat{b} estimated from SPI percentile data (99 percentile points per year, yielding ~ 30 – 35 density bins within the £25,000–£100,000 estimation range, of which ~ 5 – 8 fall within the £45,000–£55,000 exclusion window). Counterfactual density fitted with degree-7 polynomial on bins outside the exclusion window. Bootstrap standard errors from 500 residual replications. Tax year 2008/09 excluded (SPI data gap). The large 2005 SE (0.173) reflects fewer SPI percentile points in the analysis window in early years. [†]2012/13 is partially treated (HICBC effective January 2013); classified as pre-HICBC because 9 of 12 months precede the policy. Pooled means are robust to reclassifying 2012 as post-HICBC. ^aPooled SEs in brackets are cross-year standard errors of the mean (SD/\sqrt{N}): pre $SD = 0.10$, $N = 7$; post $SD = 0.12$, $N = 10$. Difference SE computed as $\sqrt{SE_{\text{pre}}^2 + SE_{\text{post}}^2}$.

The year-by-year event study is shown in [Figure 2](#). The estimates fluctuate around zero with no visible break at 2013. If anything, the post-HICBC estimates drift upward over time, from negative values in 2013–2016 to positive values in 2018–2022. This gradual drift is consistent with either slow behavioral adjustment or composition effects as nominal income growth moves more taxpayers into the £50,000 region. Among the positive post-HICBC estimates, only 2021 reaches marginal significance ($\hat{b} = 0.128$, $p = 0.029$). The statistically significant post-HICBC estimates are all negative (2014, 2015, 2016, with $p < 0.01$), suggesting missing mass rather than bunching.

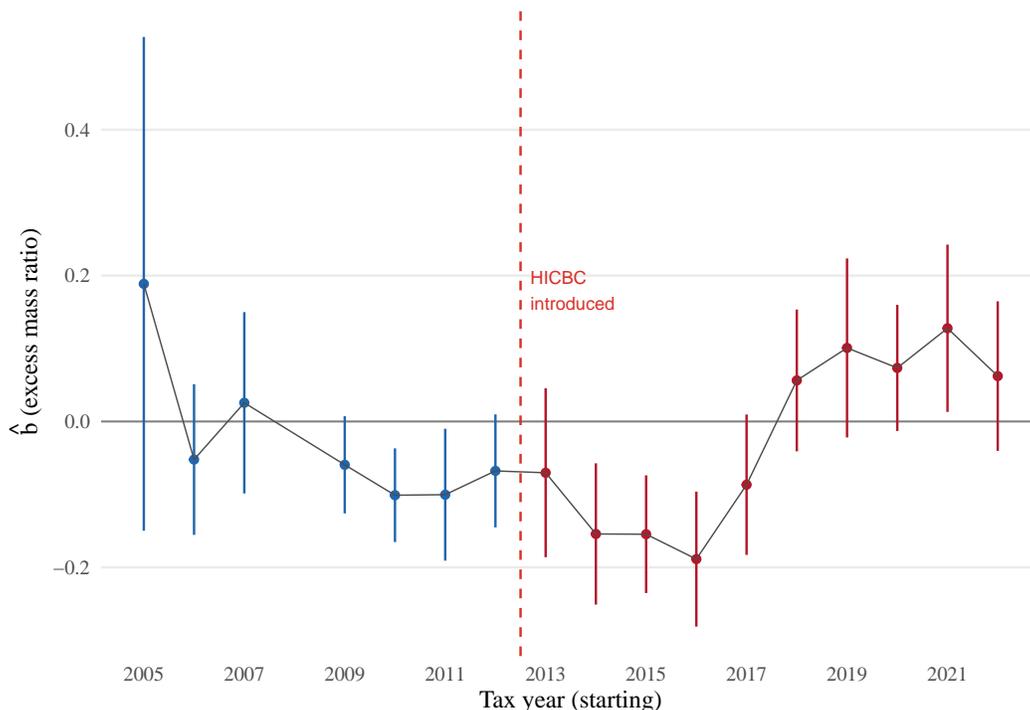


Figure 2: Year-by-Year Bunching Estimates (Event Study)

Notes: Annual excess mass ratios \hat{b} from SPI data. Vertical dashed line marks HICBC introduction (2013). Points with 95% confidence intervals from residual bootstrap. Blue = pre-HICBC; red = post-HICBC.

We do not detect bunching in total income despite a quasi-notch that, for a two-child family, implies an effective marginal rate exceeding 60 percent over a £10,000 range. As discussed in [Section 5](#), this null is consistent with treatment dilution, running variable mismatch, and limited power, as well as with the hypothesis that behavioral responses operate through channels that are invisible to this analysis. The next section investigates the available evidence on mechanisms.

6.2 Channel Decomposition

Table 3 and Figure 3 present the channel decomposition. The ASHE-based bunching estimate (PAYE employees only) shows a similar null: \hat{b}^{ASHE} averages -0.075 pre-HICBC and -0.100 post-HICBC (2013–2022), with a difference of -0.024 that is not statistically significant. However, the post-HICBC ASHE mean is dominated by a single extreme outlier in 2021 ($\hat{b} = -0.49$), driven by the furlough scheme’s distortion of PAYE employment patterns; excluding 2021, the ASHE post mean is -0.057 (SD = 0.059). The residual $\hat{b}^{\text{SPI}} - \hat{b}^{\text{ASHE}}$ (computed year-by-year for overlapping years) averages -0.007 pre-HICBC and 0.076 post-HICBC, a difference of 0.083.

Table 3: Channel Decomposition: SPI vs. ASHE Bunching

	Pre-HICBC	Post-HICBC	Difference
<i>Panel A: All taxpayers (SPI, pre: 2005–2012, post: 2013–2022)</i>			
Mean \hat{b}	-0.024	-0.023	0.001
	[0.039]	[0.038]	[0.054]
N (years)	7	10	
<i>Panel B: PAYE employees (ASHE, pre: 2009–2012, post: 2013–2022)</i>			
Mean \hat{b}	-0.075	-0.100	-0.025
	[0.005]	[0.046]	[0.047]
N (years)	4	10	
<i>Panel C: Residual (SPI – ASHE, overlapping: 2009–2012 / 2013–2022)</i>			
Mean \hat{b}	-0.007	0.076	0.083
	[0.008]	[0.064]	[0.064]
N (overlapping years)	4	10	

Notes: All SEs in brackets are cross-year standard errors of the mean (SD/ \sqrt{N}).

Panel A: SPI, pre SD = 0.103, post SD = 0.120; difference SE = $\sqrt{0.039^2 + 0.038^2}$.

Panel B: ASHE (PAYE only), degree-3 polynomial; pre SD = 0.010 ($N = 4$, years 2009–2012), post SD = 0.147 ($N = 10$). The 2021 ASHE outlier ($\hat{b} \approx -0.49$, COVID furlough artifact) dominates post variance; excluding 2021, ASHE post mean = -0.057 (SE = 0.020).

Panel C: Residual computed year-by-year for overlapping years only. The post-HICBC residual difference (0.083) is marginally significant ($t = 0.083/0.064 = 1.30$, $p = 0.19$). Tax year 2008/09 excluded from SPI.

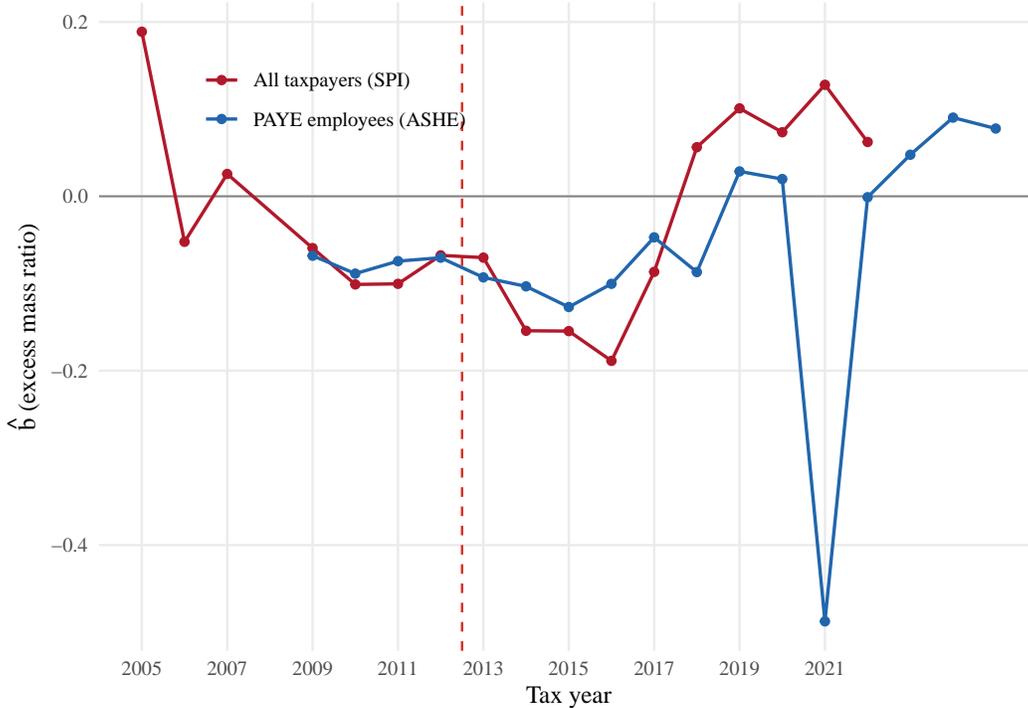


Figure 3: SPI vs. ASHE Bunching Over Time

Notes: Annual excess mass ratios from SPI (all taxpayers, red) and ASHE (PAYE employees, blue). Vertical dashed line marks HICBC introduction. The ASHE estimate for 2021 ($\hat{b} \approx -0.49$) is an extreme outlier driven by COVID-19 disruptions to PAYE employment patterns, particularly the furlough scheme’s distortion of the earnings distribution near £50,000. This value is included for completeness but should be interpreted with caution.

The interpretation of the residual is necessarily suggestive rather than definitive. The SPI and ASHE differ in population coverage (all taxpayers versus PAYE employees only), income concept (total income before tax versus annual gross earnings), data granularity (99 versus 11 percentile points), and sampling design. Computing a difference between estimates from these two distinct datasets and interpreting it as a “channel” requires strong assumptions about their comparability. Additionally, the ASHE estimates are heavily influenced by the 2021 COVID outlier, and the residual’s pre-reform period spans only four years (2009–2012), limiting the precision of the pre–post comparison. This comparison is therefore consistent with, but does not establish, a small non-PAYE income manipulation response among self-employed taxpayers after the HICBC. The magnitude is small compared to the administrative response, and the statistical uncertainty ($p = 0.19$) does not permit strong inference.

6.3 Administrative Evidence

Figure 4 documents the administrative response. The number of families opting out of Child Benefit grew from 43,000 in 2013 to 740,000 in 2023, before declining slightly to 712,000 after the 2024 threshold increase. HICBC revenue grew from £416 million (2018) to £525 million (2022), and the number of liable individuals from 293,000 to 440,000 over the same period.

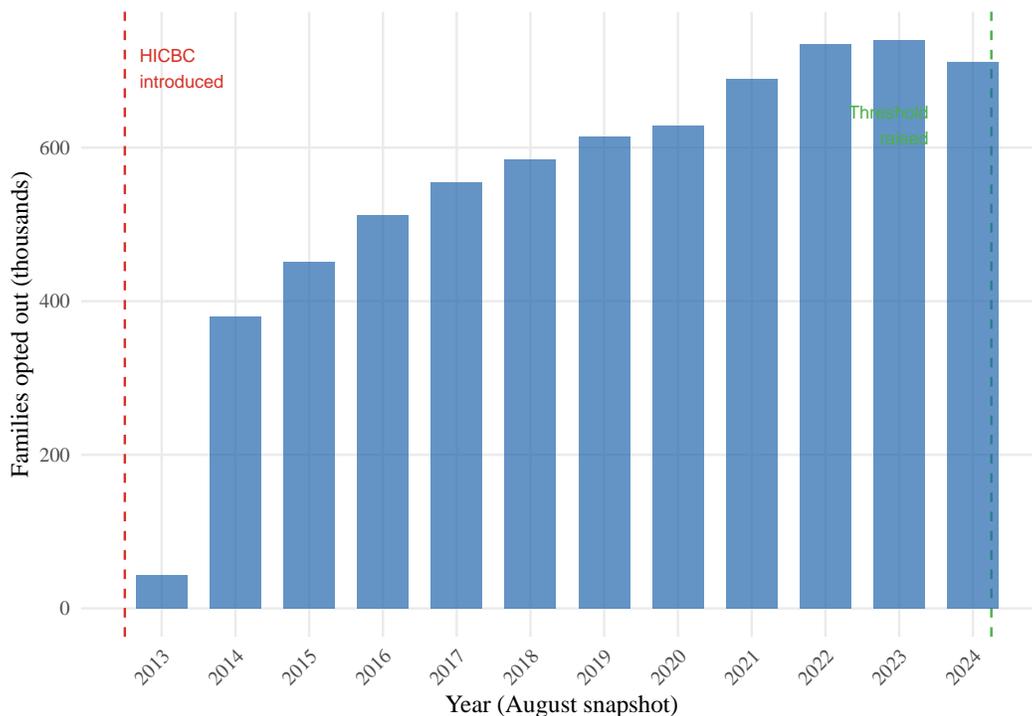


Figure 4: Families Opting Out of Child Benefit, 2013–2024

Notes: Number of families electing not to receive Child Benefit (thousands). August snapshot each year. Red dashed line marks HICBC introduction; green dashed line marks 2024 threshold increase to £60,000. Source: HMRC Child Benefit statistics.

The combined numbers are striking. By 2022/23, 735,000 families had opted out of receiving Child Benefit entirely, while a further 440,000 individuals were liable for the HICBC charge through self-assessment. Together, the charge had a direct administrative impact on a substantial share of Child Benefit-eligible families. Yet none of this massive behavioral disruption is visible in the income distribution.

6.4 Robustness

The null bunching result is robust across a wide range of specifications. I examine sensitivity to four dimensions: polynomial degree, exclusion window width, round-number placebos, and

pre-reform placebo timing.

Polynomial degree sensitivity is tested using degrees 5, 7, 9, and 11. The mean post-HICBC bunching estimate ranges from -0.023 (degree 7, the baseline) to $+0.001$ (degree 5), with all estimates statistically indistinguishable from zero. The standard deviations across years are similar (0.10–0.12) regardless of polynomial degree, confirming that the null is not an artifact of counterfactual specification (see [Table 4](#) and [Figure 6](#)).

Exclusion window sensitivity is tested using windows of $\pm£3,000$, $\pm£5,000$ (baseline), $\pm£7,000$, and $\pm£10,000$. Mean estimates range from -0.008 to $+0.044$, all near zero and statistically insignificant ([Table 5](#) and [Figure 7](#)). The wider window ($\pm£10,000$) shows a modestly positive mean of 0.044, possibly reflecting that a broader window captures more of the HICBC withdrawal range, but it remains well within standard error bounds.

Round-number placebo tests compare bunching at $£40,000$, $£45,000$, $£50,000$, $£55,000$, and $£60,000$ ([Figure 5](#)). The mean post-HICBC \hat{b} at $£50,000$ (-0.003 , from the placebo procedure which uses a wider estimation range of $£20,000$ – $£100,000$) is actually smaller in absolute value than the estimates at $£40,000$ (0.004) and $£45,000$ (0.017), confirming that there is nothing special about the income distribution at the HICBC threshold. The difference between this estimate (-0.003) and the main specification (-0.023) reflects the sensitivity of the polynomial counterfactual to the estimation range; both values are firmly in the null range. The estimate at $£60,000$ (0.463) is large but extremely noisy ($SD = 1.076$), driven by the sparseness of the density at higher incomes.

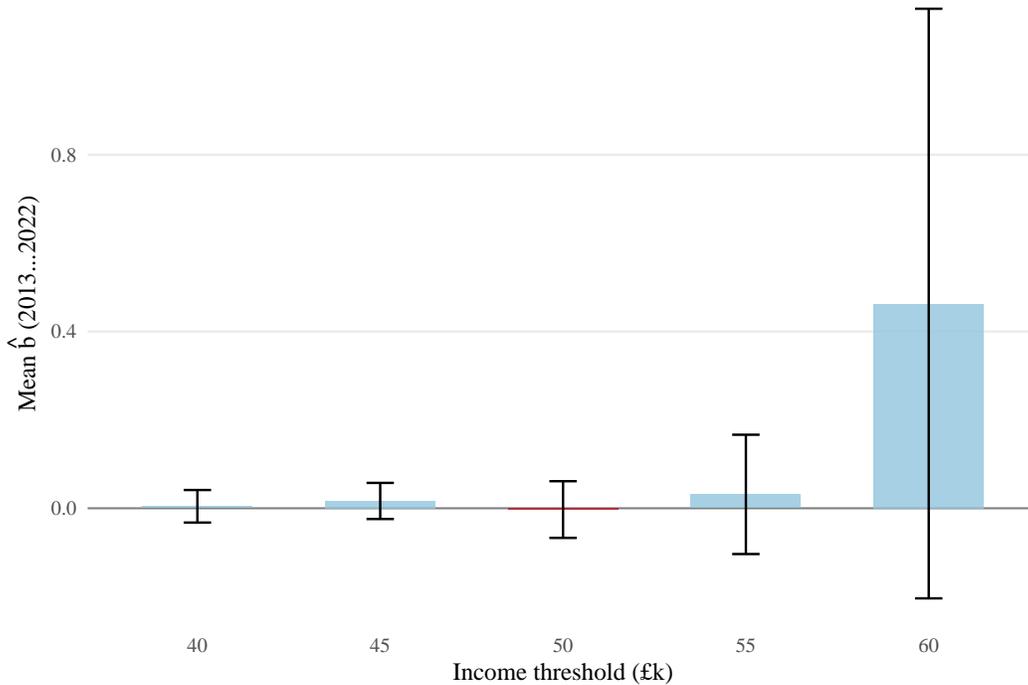


Figure 5: Placebo Tests at Round-Number Income Thresholds

Notes: Mean post-HICBC (2013–2022) excess mass ratio at each round-number threshold. Red bar indicates the actual HICBC threshold (£50,000). Error bars show 95% confidence intervals.

Finally, the pre-reform placebo confirms that the null is not driven by pre-existing bunching. The mean pre-2013 bunching estimate is -0.024 with a standard deviation of 0.10 , centered on zero. Individual pre-reform years range from -0.10 to $+0.19$, with the positive outlier (2005) reflecting noise in a single year rather than systematic bunching.

7. Discussion

7.1 Why No Bunching? Three Mechanisms

The absence of bunching in total income, despite a powerful notch and massive administrative response, can be explained by three complementary mechanisms.

Mechanism 1: The pension contribution channel (hypothesized). Among taxpayers earning £50,000–£60,000, over 62 percent already have pension contributions, with mean relief of £8,634 (Table 1, Panel D). For many of these taxpayers, increasing pension contributions by a few thousand pounds would bring their ANI below the threshold, eliminating the HICBC charge. Since pension contributions reduce ANI but not total income,

this response would be invisible in the SPI distribution. The tax advantage of pension relief (higher-rate relief at 40 percent plus avoiding the HICBC) makes this an attractive option in principle—taxpayers who pursue it effectively receive double tax relief on their pension saving.

This mechanism is consistent with the pattern of financial adviser recommendations that emerged after 2013. HMRC’s own documentation notes that salary sacrifice arrangements are “the most tax-efficient way” for employees to reduce ANI below the threshold, and the financial planning industry quickly developed products targeting HICBC-affected families. However, because the SPI records total income rather than ANI, we cannot directly observe this channel. The evidence is consistent with pension avoidance being quantitatively important, but cannot establish it.

Mechanism 2: Administrative exit. The 712,000 families who opted out of Child Benefit represent a clean administrative response: rather than adjusting income, they simply stopped receiving the benefit. For a one-child family receiving £1,133 per year, the benefit of claiming may not justify the hassle of filing self-assessment and paying the charge, particularly for higher earners well above £60,000. The opt-out rate grew steadily from 2013 to 2023, suggesting gradual awareness rather than immediate response.

Mechanism 3: Limited real labor supply response. The frictions documented by [Chetty et al. \(2011\)](#)—wage rigidity, hours constraints, job change costs—make real income adjustment the most expensive margin. For PAYE employees, annual salary is typically set by contract and not easily adjusted in increments of a few thousand pounds. The self-employed have more flexibility, and the modest positive shift in the SPI-ASHE residual is consistent with some self-employment income manipulation, but the magnitude is small.

These three mechanisms form a hierarchy of adjustment costs, from cheapest (opt-out) to most expensive (real income reduction). Standard bunching analysis, by examining only total income, captures only the most expensive channel. This is analogous to the distinction between the elasticity of taxable income and the elasticity of broad income emphasized by [Kopczuk \(2005\)](#) and [Chetty \(2009\)](#): the welfare-relevant elasticity depends on which margins are activated.

7.2 Implications for Bunching as a Method

The HICBC case carries a cautionary message for the bunching literature. Bunching methods have been enormously productive in estimating behavioral responses to taxation ([Kleven, 2016](#)). But they measure bunching in a *specific income variable*, and this may diverge from the policy-relevant income concept when deductions, exemptions, or administrative avoidance provide pathways between the two.

In the canonical applications—the Pakistani income tax notches studied by [Kleven and Waseem \(2013\)](#), the US EITC kinks studied by [Saez \(2010\)](#), Danish tax kinks studied by [Chetty et al. \(2011\)](#)—the income concept in the data aligns closely with the income concept in the policy rule. The HICBC is different: the notch applies to ANI, but the observable distribution is total income. This wedge is created by pension contributions, which are neither illegal nor distortionary—indeed, they represent exactly the kind of welfare-improving tax planning that a well-designed system should encourage.

The implication is not that bunching methods are unreliable, but that they must be interpreted with care regarding the adjustment margin ([Chetty, 2012](#)). A null bunching result does not mean “no behavioral response”—it means “no response on the specific margin measured.” When cheaper margins exist, bunching may understate the total behavioral elasticity ([Saez et al., 2012](#)).

7.3 The Time-Series Pattern

The event study ([Figure 2](#)) reveals a suggestive time pattern that merits careful interpretation. In the first four years after the HICBC (2013–2016), bunching estimates are negative (−0.07 to −0.19), similar to or below the pre-reform baseline. From 2017 onward, estimates shift upward (−0.09 to +0.13), with three of the five later years showing positive point estimates, though none of the positive estimates reaches statistical significance individually.

This gradual shift admits at least three interpretations. First, it may reflect genuine behavioral adjustment that takes time to emerge as taxpayers learn about the HICBC, seek financial advice, and restructure their affairs. The bunching literature has documented such learning effects in other contexts: [Chetty et al. \(2013\)](#) find that EITC bunching is stronger in neighborhoods with more prior exposure, and [Mortenson and Whitten \(2020\)](#) document that bunching at US tax kinks builds over time. If taxpayers are slowly learning to adjust their total income (not just their ANI), the positive estimates after 2017 may represent the beginning of a real response.

Second, the pattern may be compositional. As nominal incomes rise, different taxpayers populate the £50,000 region in each year. Those at £50,000 in 2022 occupy a different position in the ability distribution than those at £50,000 in 2013. If higher-ability taxpayers are better at adjusting their incomes, the composition shift could generate apparent bunching even without behavioral change.

Third, the positive estimates in 2018–2022 may reflect noise. The standard errors (≈ 0.04 – 0.06) are large enough that the positive estimates are consistent with a true effect of zero. With 10 post-reform years, several would be expected to deviate from zero by chance. The pre-reform period contains a similarly large positive outlier (2005: $\hat{b} = 0.19$) that clearly

predates the HICBC.

Disentangling these interpretations would require individual-level panel data linking tax records over time—precisely the kind of administrative microdata that the UK’s ADR UK RAPID program aims to make available for research. With individual panels, one could track whether specific taxpayers adjust their income downward as they approach the threshold, controlling for individual fixed effects and lifecycle income trajectories. The quantile-based approach used here, while valuable for detecting aggregate responses, cannot resolve individual-level dynamics.

7.4 Limitations

This analysis has three principal limitations that bear on how the results should be interpreted.

Treatment dilution. The SPI covers all UK income taxpayers, but HICBC applies only to those actively claiming Child Benefit—roughly 13 percent of taxpayers near £50,000. As shown in [Section 5](#), this dilution means that moderate behavioral responses among the treated subpopulation would be statistically undetectable in the all-taxpayer distribution. Our confidence intervals are too wide to rule out behavioral bunching estimates up to $b_{\text{treated}} \approx 0.50$ among the genuinely affected population. Addressing this limitation requires individual-level microdata linking Child Benefit claim status to income records—data that are available within the UK’s administrative system but not publicly released at fine granularity.

Running variable mismatch. The HICBC is assessed on adjusted net income (ANI), but the SPI records total income before tax. Pension contributions, the primary avoidance mechanism, reduce ANI without changing total income. Any behavioral bunching in ANI is therefore invisible to this analysis. The mismatch is not merely an empirical inconvenience—it represents the core mechanism through which many families likely adjust. The absence of bunching in total income should not be interpreted as an absence of bunching in the policy-relevant variable.

Limited power and data granularity. The 99 percentile points in the SPI yield approximately 5–8 density bins within the £45,000–£55,000 exclusion window, and the all-taxpayer dilution problem compounds this. As discussed in [Section 5](#), the standard errors on annual bunching estimates (≈ 0.04 – 0.06) are large enough that only very large responses among affected families would be detectable. The 10 post-reform years of data are insufficient to substantially reduce cross-year variance through averaging. A formal power analysis suggests that the analysis would require either individual-level microdata or approximately 40 post-reform years to detect a realistic treated-population response of $b_{\text{treated}} = 0.20$ with 80 percent power.

These limitations do not invalidate the paper’s contribution—documenting the divergence

between bunching nulls and administrative responses remains informative—but they do counsel against interpreting the null as evidence of no behavioral response among HICBC-affected families (Chetty, 2012).

7.5 Welfare Implications

The welfare analysis of the HICBC depends critically on which adjustment margin dominates. Following the framework of Chetty (2009), the deadweight loss of a tax depends not on the overall elasticity of taxable income but on the elasticity of “broad income”—income measured before deductions and avoidance opportunities. If behavioral responses operate through deductions that are themselves efficient (like pension saving), the fiscal externality is limited and the deadweight loss is smaller than the total revenue loss would suggest (Slemrod, 1996; Saez et al., 2012).

To illustrate the welfare implications *conditional on the pension mechanism hypothesis*, consider a taxpayer who increases pension contributions by £5,000 to bring their ANI below £50,000. The HICBC charge avoided is approximately £1,885 (for a two-child family). The “cost” of avoidance is the liquidity cost of pension contributions—the money is locked until retirement but earns a return. At higher-rate tax relief (40 percent), the net cost of a £5,000 pension contribution is approximately £3,000 (after tax relief). The taxpayer gains £1,885 in retained Child Benefit and the long-run return on £5,000 of additional pension savings. Under this interpretation, the HICBC would effectively subsidize retirement saving for those who use the pension channel. This analysis is illustrative: we cannot directly observe whether families are using this mechanism, and the welfare implications differ substantially under alternative interpretations.

The welfare calculus is less favorable for the 712,000 families who opted out of Child Benefit entirely. These families forgo a benefit they are legally entitled to, either because they are unaware of the pension avoidance strategy, because they cannot easily increase pension contributions (e.g., the self-employed), or because the administrative hassle of filing self-assessment outweighs the benefit amount. For a one-child family earning £70,000, the optimal strategy is to opt out (the full benefit would be clawed back anyway), and the welfare cost is the lost benefit of £1,133 per year. For a two-child family earning £52,000, the optimal strategy is to increase pension contributions by £2,000—but if the family lacks financial literacy or employer-provided salary sacrifice, they may instead opt out, losing £1,885 annually.

The distributional pattern is therefore regressive within the affected population: financially sophisticated families with access to employer pension schemes exploit the pension channel, retaining both the Child Benefit and a tax-advantaged pension contribution. Less sophisticated

families or those without employer pension access either pay the charge or opt out. The HICBC thus creates a wedge between the “savvy” and the “naive” that is not captured by standard distributional analysis based on income percentiles.

7.6 Policy Implications

The HICBC’s design creates perverse incentives that primarily penalize families who are unaware of avoidance strategies or lack access to them. Financially sophisticated families use pension contributions to avoid the charge entirely while building retirement savings. Less sophisticated families either pay the charge (generating £525 million in revenue) or opt out of a benefit they are entitled to. The distributional consequences are regressive within the affected income range: the charge falls most heavily on those least equipped to avoid it.

The 2024 reform raising the threshold to £60,000 addresses the most acute problems but does not resolve the underlying design flaw. The HICBC still applies to individual rather than household income, creating incentives for families to arrange their affairs so that neither partner exceeds the threshold. And the withdrawal range (£60,000–£80,000) still creates elevated marginal rates, albeit lower than before.

The broader lesson is that benefit withdrawal design should account for all available adjustment margins, not just the one that policymakers intend to influence. When administrative exit is easy (as with Child Benefit opt-outs) or when the tax base can be narrowed through deductions (as with pension contributions), the effective notch in real economic activity is much smaller than the statutory notch in the tax code. This echoes the recommendation of the Mirrlees Review ([Mirrlees et al., 2011](#)) for a more integrated approach to tax and benefit design—one that recognizes that taxpayers optimize across all available margins, and that the administrative complexity of a system determines which margins are cheapest to exploit.

International comparison underscores this point. Benefit withdrawal systems in other countries—the US EITC phase-out, the Australian Medicare Levy Surcharge, the Canadian Canada Child Benefit phase-out—all create implicit marginal tax rates in specific income ranges. But the HICBC is unusual in the steepness of the taper (100 percent withdrawal over a £10,000 range, creating a quasi-notch) combined with the availability of a low-cost avoidance strategy (pension contributions). Most benefit withdrawal systems either phase out more gradually or apply to income concepts that are harder to manipulate. The HICBC’s combination of a steep quasi-notch with an accessible escape route makes it a particularly instructive case for understanding how the design of avoidance opportunities shapes behavioral responses to tax incentives.

7.7 Comparison to Other Bunching Studies

The null bunching result places this paper in a distinct position within the bunching literature. Most studies that apply bunching methods to tax notches find detectable excess mass, though the magnitudes vary widely. [Kleven and Waseem \(2013\)](#) find substantial bunching at Pakistani income tax notches, with excess mass ratios of 2–4 for the self-employed but near zero for wage earners—a pattern consistent with the role of third-party reporting in limiting responses ([Kleven et al., 2011](#)). [Best et al. \(2015\)](#) find significant bunching at firm-level revenue notches in Pakistan. [Almunia and Lopez-Rodriguez \(2018\)](#) document bunching at monitoring thresholds in Spain.

In developed-country settings, bunching tends to be smaller. [Bastani and Selin \(2014\)](#) find modest bunching at Swedish tax kinks. [Devereux et al. \(2014\)](#) find limited bunching in UK corporate taxable income. [Le Maire and Schjerning \(2013\)](#) find that bunching in Denmark is concentrated among the self-employed. The common thread is that bunching is largest where adjustment frictions are lowest—self-employment income, small firms, cash-intensive businesses—and smallest where third-party reporting constrains manipulation.

The HICBC case extends this finding in an important direction: even when the notch is powerful and the affected population is economically sophisticated (upper-middle-income parents), bunching may be absent if the cheapest adjustment channel operates through a dimension that is not captured in the bunching analysis. This suggests that the bunching literature’s focus on income distributions may systematically understate behavioral responses in settings where tax-base-narrowing deductions provide an alternative to income adjustment.

8. Conclusion

The UK’s High Income Child Benefit Charge created one of the steepest benefit tapers in any developed country’s tax-benefit system, generating effective marginal rates exceeding 60 percent over a £10,000 range. Standard bunching theory predicts detectable income clustering just below the £50,000 threshold. No such pattern appears in published total income distributions. Yet the administrative response has been massive: over 700,000 families opted out of Child Benefit, take-up fell by 9 percentage points, and HMRC collects over half a billion pounds annually from those who continue to claim.

This divergence between statistical bunching and administrative response reveals a hierarchy of adjustment margins that standard bunching methods cannot fully capture. When the cheapest responses are administrative (opting out) or operate through deductions (pension contributions that reduce adjusted net income without changing total income), the income distribution—the object that bunching methods examine—is the wrong place to look for

behavioral responses.

The practical implication is straightforward: bunching estimates bound the response on the measured income margin, but the total behavioral response may be much larger. In the HICBC case, the bound in total income is approximately zero, while the true response involves over a million families changing their relationship with a universal benefit. Researchers applying bunching methods should identify all adjustment margins available to taxpayers and assess whether the income concept in their data captures the relevant ones. When it does not, the bunching estimate measures a floor, not a ceiling, on behavioral response. Treatment dilution—the fact that the policy applies to a subset of taxpayers who cannot be isolated in the aggregate distribution—can compound this problem further.

These results should be interpreted with the caveat that the analysis uses published total income distributions for all taxpayers, not the treated subpopulation or the policy-relevant income concept (ANI). Subject to these data limitations, the evidence is consistent with behavioral responses to the HICBC operating primarily through administrative channels and pension deductions rather than through adjustments to total income—a pattern that, if confirmed with microdata, would carry important implications for the welfare costs of benefit withdrawal design.

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

- Adam, Stuart and Helen Miller**, “Taxable Income Responses to Tax Changes: Evidence from the UK,” *Fiscal Studies*, 2021, 42 (2), 265–299.
- Almunia, Miguel and David Lopez-Rodriguez**, “Under the Radar: The Effects of Monitoring Firms on Tax Compliance,” *American Economic Journal: Economic Policy*, 2018, 10 (1), 1–38.
- Bastani, Spencer and Håkan Selin**, “Bunching and Non-Bunching at Kink Points of the Swedish Tax Schedule,” *Journal of Public Economics*, 2014, 109, 36–49.
- Best, Michael Carlos, Anne Brockmeyer, Henrik Jacobsen Kleven, Johannes Spinnewijn, and Mazhar Waseem**, “Production versus Revenue Efficiency with Limited Tax Capacity: Theory and Evidence from Pakistan,” *Journal of Political Economy*, 2015, 123 (6), 1311–1355.
- Brewer, Mike, Emmanuel Saez, and Andrew Shephard**, “Means-Testing and Tax Rates on Earnings,” *Dimensions of Tax Design: The Mirrlees Review*, 2010, pp. 90–173.
- Browne, James and Andrew Hood**, “A Survey of the UK Tax System,” *IFS Briefing Note BN09*, 2012.
- Chetty, Raj**, “Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance,” *American Economic Journal: Economic Policy*, 2009, 1 (2), 31–52.
- , “Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply,” *Econometrica*, 2012, 80 (3), 969–1018.
- , **John N. Friedman, and Emmanuel Saez**, “Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings,” *American Economic Review*, 2013, 103 (7), 2683–2721.
- , – , **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.
- Devereux, Michael P., Li Liu, and Simon Loretz**, “The Elasticity of Corporate Taxable Income: New Evidence from UK Tax Records,” *American Economic Journal: Economic Policy*, 2014, 6 (2), 19–53.

- Feldstein, Martin**, “Tax Avoidance and the Deadweight Loss of the Income Tax,” *Review of Economics and Statistics*, 1999, 81 (4), 674–680.
- Gruber, Jonathan and Emmanuel Saez**, “The Elasticity of Taxable Income: Evidence and Implications,” *Journal of Public Economics*, 2002, 84 (1), 1–32.
- HM Revenue and Customs**, “Child Benefit, Child Tax Credit and Working Tax Credit: Take-up Rates 2021 to 2022,” Technical Report, HMRC 2023.
- , “Survey of Personal Incomes: Distribution of Individual Total Income,” Technical Report, HMRC 2024. Tables 3.1–3.17.
- Kleven, Henrik Jacobsen**, “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.
- , **Martin B. Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez**, “Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark,” *Econometrica*, 2011, 79 (3), 651–692.
- Kopczuk, Wojciech**, “Tax Bases, Tax Rates and the Elasticity of Reported Income,” *Journal of Public Economics*, 2005, 89 (11–12), 2093–2119.
- Maire, Daniel Le and Bertel Schjerning**, “Tax Bunching, Income Shifting and Self-Employment,” *Journal of Public Economics*, 2013, 107, 1–18.
- Mirrlees, James A., Stuart Adam, Timothy Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, and James M. Poterba**, “Tax by Design: The Mirrlees Review,” *Institute for Fiscal Studies*, 2011.
- Mortenson, Jacob A. and Andrew Whitten**, “How Sensitive Are Taxpayers to Marginal Tax Rates? Evidence from Income Bunching in the United States,” *Journal of Public Economics*, 2020, 183, 104136.
- Office for Budget Responsibility**, “Economic and Fiscal Outlook, Spring Budget 2024,” Technical Report, OBR 2024.
- Office for National Statistics**, “Annual Survey of Hours and Earnings: Percentile Data,” Technical Report, ONS 2024. NOMIS Table 1.

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

– , **Joel Slemrod, and Seth H. Giertz**, “The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review,” *Journal of Economic Literature*, 2012, 50 (1), 3–50.

Slemrod, Joel, “High-Income Families and the Tax Changes of the 1980s: The Anatomy of Behavioral Response,” *NBER Book Chapter*, 1996, pp. 169–189.

A. Data Appendix

A.1 HMRC Survey of Personal Incomes

The SPI percentile data is sourced from HMRC Table 3.1a, which reports the 1st through 99th percentile points of total income before tax for the United Kingdom. The data is available from 1999/2000 to 2022/23. I download the ODS file directly from GOV.UK.⁴

Income concept: Total income before tax includes employment income, self-employment profits, pension income, investment income, and other income. It does *not* subtract pension contributions, Gift Aid, or other reliefs. This is the key distinction from adjusted net income, which does subtract these items.

Sample construction: All UK income taxpayers are included. The SPI is based on a stratified sample of approximately 700,000 tax records, reweighted to the population of approximately 31–35 million taxpayers depending on year.

Density construction: From 99 percentile points, I compute bin densities as $\hat{f}(z) = (p_{j+1} - p_j) / (100 \times (z_{j+1} - z_j))$ where p_j and p_{j+1} are adjacent percentile indices and z_j, z_{j+1} are the corresponding income levels. I add $P_0 = \text{£}0$ and $P_{100} = 1.5 \times P_{99}$ as boundary conditions. This yields approximately 100 density bins per year, with bin widths of $\text{£}1,500\text{--}\text{£}3,000$ in the $\text{£}30,000\text{--}\text{£}80,000$ range.

A.2 Annual Survey of Hours and Earnings

ASHE data is accessed through the NOMIS API (<https://www.nomisweb.co.uk/>). I retrieve annual gross earnings percentiles (P10, P20, P25, P30, P40, P50, P60, P70, P75, P80, P90) for full-time employees, all occupations, UK-wide, from 2002 to 2024. Separate requests retrieve male and female breakdowns.

Coverage: ASHE covers employees on PAYE, based on a 1% sample of National Insurance numbers. It excludes the self-employed, company directors paying themselves through dividends, and those not in PAYE employment. This makes it a natural comparator for the all-taxpayer SPI. The pre-HICBC ASHE sample is limited to 2009–2012 ($N = 4$) because the $\text{£}50,000$ threshold falls within the 80th–90th percentile bin only from 2009 onward; in earlier years, the bin boundaries place $\text{£}50,000$ in the 90th–100th percentile bin, which is too wide for meaningful density estimation.

Density construction: With 11 percentile points, ASHE density bins are coarser than SPI bins. I add $P_0 = 0$ and $P_{100} = 2 \times P_{90}$ as boundaries. I use a degree-3 polynomial for ASHE counterfactual fitting (versus degree-7 for SPI) to avoid overfitting with fewer data

⁴HMRC, “Percentile points from 1 to 99 for total income before and after tax,” available at gov.uk.

points.

A.3 Child Benefit Administrative Data

Administrative statistics are compiled from published HMRC releases on Child Benefit, HICBC, and benefit take-up. Sources include:

- Child Benefit claimant statistics (annual August snapshot)
- HICBC self-assessment data (annual)
- Child Benefit, Child Tax Credit and Working Tax Credit take-up statistics

All figures are as published by HMRC with no further processing.

B. Identification Appendix

B.1 Pre-Reform Placebo

If the HICBC generates bunching, bunching estimates should be approximately zero in the pre-reform period (2005–2012). The pre-HICBC panel of [Table 2](#) confirms this: the mean pre-2013 \hat{b} is -0.024 (SD = 0.103), with individual years ranging from -0.101 (2010) to $+0.189$ (2005). The wide range reflects estimation noise from the quantile-based density, not systematic bunching.

C. Robustness Appendix

C.1 Polynomial Degree Sensitivity

Table 4: Sensitivity to Polynomial Degree

Polynomial Degree	Mean \hat{b} (2013–2022)	SE	SD
5	0.001	[0.032]	0.102
7	-0.023	[0.038]	0.120
9	-0.014	[0.031]	0.097
11	0.000	[0.039]	0.123

Notes: Mean excess mass ratio across post-HICBC years (2013–2022) for different polynomial degrees used in counterfactual density estimation. SE in brackets is the cross-year standard error of the mean ($SD/\sqrt{10}$). SD is across years.

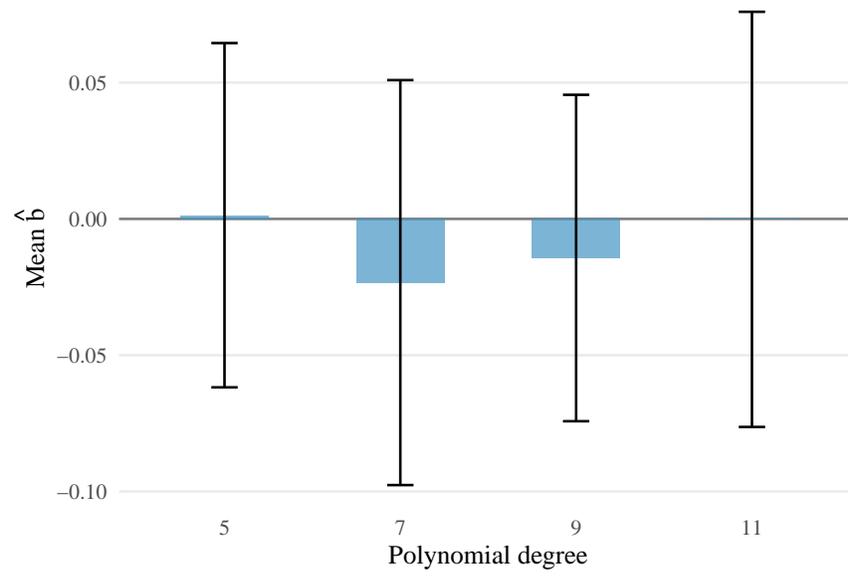


Figure 6: Polynomial Degree Sensitivity

C.2 Exclusion Window Sensitivity

Table 5: Sensitivity to Exclusion Window (Non-Baseline Estimation Range: £20k–£100k)

Window	Mean \hat{b} (2013–2022)	SE	SD
±£3,000	0.014	[0.034]	0.106
±£5,000	−0.003	[0.033]	0.104
±£7,000	−0.008	[0.035]	0.110
±£10,000	0.044	[0.015]	0.047

Notes: Mean excess mass ratio across post-HICBC years for different exclusion window widths around £50,000. Baseline is ±£5,000. Degree-7 polynomial. SE in brackets is the cross-year standard error of the mean ($SD/\sqrt{10}$). These estimates use an estimation range of £20,000–£100,000 (versus £25,000–£100,000 in the main specification), which accounts for the difference between the ±£5,000 estimate here (−0.003) and the main result in [Table 2](#) (−0.023). Both are well within the null range; the sensitivity to estimation range confirms that the precise point estimate depends on the polynomial’s extrapolation from distant bins, while the qualitative conclusion of no significant bunching is invariant.

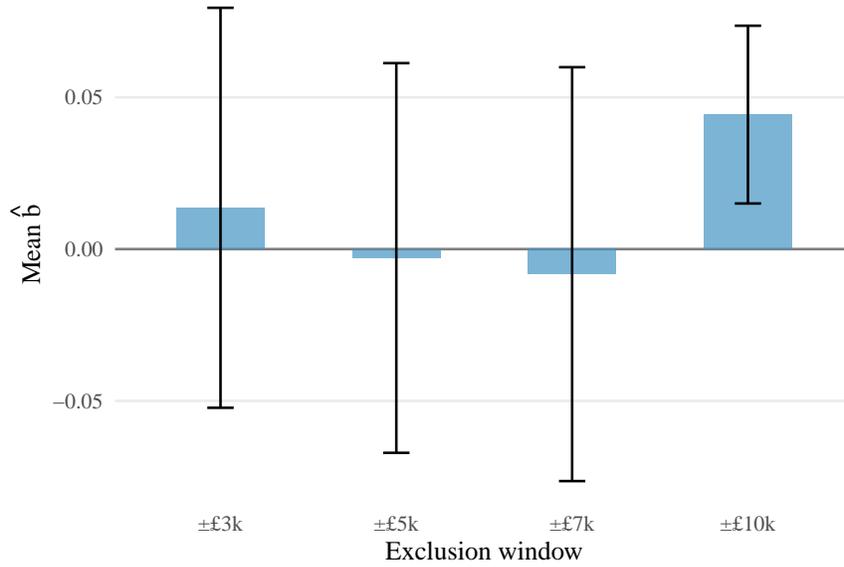


Figure 7: Exclusion Window Sensitivity

C.3 Round-Number Placebo Estimates

Table 6: Placebo Bunching at Round-Number Thresholds (Non-Baseline Estimation Range: £20k–£100k)

Threshold	Mean \hat{b} (2013–2022)	SE	SD	N
£40,000	0.004	[0.019]	0.059	10
£45,000	0.017	[0.021]	0.066	10
£50,000	-0.003	[0.033]	0.104	10
£55,000	0.031	[0.069]	0.218	10
£60,000	0.463	[0.340]	1.076	10

Notes: Excess mass ratio at each round-number threshold using degree-7 polynomial, ±£5,000 exclusion window, and £20,000–£100,000 estimation range. SE in brackets is the cross-year standard error of the mean ($SD/\sqrt{10}$). The £50,000 estimate corresponds to the HICBC threshold. The large estimate at £60,000 reflects extreme noise from the sparseness of the density at higher incomes.

C.4 Pension Contributions by Income Band

Table 7: Pension Contributions by Income Band

Income Band	Taxpayers (000s)	With Pension (000s)	% with Pension	Mean Pension (£)
£20,000	9,910	4,820	48.6	1,774
£30,000	5,990	3,270	54.6	2,887
£40,000	3,680	2,250	61.1	5,733
£50,000	3,270	2,050	62.7	8,634
£70,000	1,550	1,020	65.8	12,941
£100,000	749	441	58.9	21,429

Notes: Income band indicates lower bound of HMRC income range. Source: HMRC Table 3.5.

D. Heterogeneity Appendix

D.1 Gender Decomposition

If the HICBC response operates through salary sacrifice (available to PAYE employees), it should be visible in the earnings distribution of employees near £50,000. I decompose the ASHE density at £50,000 by gender over time.

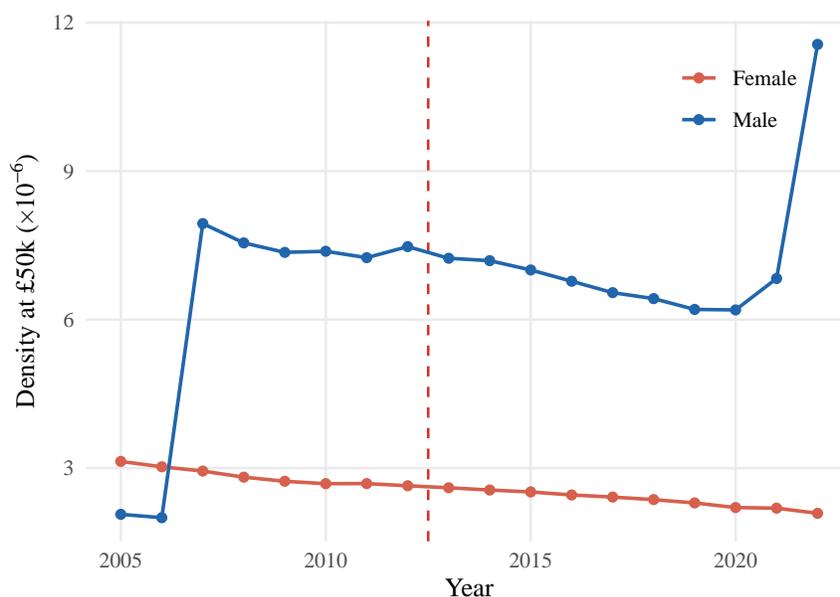


Figure 8: ASHE Density at £50,000 by Gender

Notes: Density of annual earnings at the £50,000 threshold from ASHE, separately for male and female full-time employees. Vertical line marks HICBC introduction.

D.2 Administrative Response Over Time

Table 8: Child Benefit Administrative Data

<i>Panel A: Opt-Outs and Take-up (2013–2024)</i>			
Year	Opted Out (000s)	Take-up (%)	
2013	43	97	
2015	451	93	
2017	555	91	
2019	614	90	
2021	690	89	
2023	740	88	
2024 [†]	712	n.a. [†]	

<i>Panel B: HICBC Liability and Revenue (2018–2022)</i>			
Year	HICBC Liabile (000s)	Revenue (£m)	Take-up (%)
2018	293	416	91
2019	352	468	90
2020	373	395	90
2021	393	470	89
2022	440	525	88

Notes: Panel A: Opt-out figures are August snapshots; selected years shown. [†]The 2024 take-up rate is not yet published; the decline from 740,000 to 712,000 opt-outs reflects families returning to claiming after the threshold increase to £60,000. Panel B: HICBC liable and revenue figures are available from 2018/19 onward only. Revenue for the tax year ending in the given calendar year. Source: HMRC published statistics.

E. Standardized Effect Sizes

This paper’s primary estimates are excess mass ratios (\hat{b}) from a bunching design, not regression coefficients from a treatment-effect framework. The bunching estimand is itself a standardized measure: it expresses excess density as a fraction of counterfactual density, making it comparable across settings without further normalization. Unlike a DiD or RDD coefficient that carries outcome-specific units, \hat{b} is dimensionless by construction.

Accordingly, the standard SDE table (reporting $\hat{\beta}/\text{SD}(Y)$) is not directly applicable. For comparability with the broader literature, I report the main estimates in a format that facilitates cross-study comparison.

Table 9: Summary of Bunching Estimates

Estimand	Specification	Estimate	SE	Classification
Pre-HICBC \hat{b}	SPI, degree-7, $\pm\text{£}5\text{k}$	-0.024	0.039	Null
Post-HICBC \hat{b}	SPI, degree-7, $\pm\text{£}5\text{k}$	-0.023	0.038	Null
Difference (post - pre)	SPI, degree-7, $\pm\text{£}5\text{k}$	0.001	0.054	Null

Notes: Excess mass ratios (\hat{b}) are dimensionless: they measure the percentage deviation of observed density from counterfactual density within the exclusion window around $\text{£}50,000$. Classification uses the standard thresholds: null ($|b| < 0.05$), small (0.05 to 0.10), large (> 0.10). The main result—the pre-post difference—is null (0.001). SEs are cross-year standard errors of the mean (SD/\sqrt{N}): pre-HICBC $\text{SD} = 0.10$, $N = 7$; post-HICBC $\text{SD} = 0.12$, $N = 10$. Difference $\text{SE} = \sqrt{0.039^2 + 0.038^2} = 0.054$. **Research question:** Does the UK’s High Income Child Benefit Charge (HICBC), which claws back Child Benefit for families where the higher earner exceeds $\text{£}50,000$ ANI, generate bunching in the total income distribution? **Data:** HMRC Survey of Personal Incomes (99 percentile points), 2005–2022, all UK income taxpayers (~ 31 – 35 million). **Method:** Polynomial counterfactual fitting with exclusion window; bootstrap SE for individual years, cross-year SE for means. **Sample:** Total income $\text{£}25,000$ – $\text{£}100,000$ range for density estimation.