Online Appendix to "Bunching to Maximize Tax Credits: Evidence from Kinks in the U.S. Tax Schedule"

Jacob A. Mortenson*
Joint Committee on Taxation,
U.S. Congress
502 Ford House Office Building
Washington, DC 20515
jake.mortenson@jct.gov

Andrew Whitten
Office of Tax Analysis,
U.S. Department of the Treasury
1500 Pennsylvania Avenue NW
Washington, DC 20220
andrew.whitten@treasury.gov

Appendix A: Bunching Estimation

In this appendix, we provide a more exhaustive discussion of the bunching we see in the data. We also show how our bunching estimates vary under alternative parameter choices and polynomial degrees.

A1. Bunching Estimation Results

Figure A1 depicts many of the kinks we study for an unmarried taxpayer with two children in 2014.[1] The horizontal axis is wage income, and the vertical axis is the marginal tax rate, ignoring state taxes. Kinks with increasing marginal tax rates are convex, while those with decreasing marginal tax rates are non-convex. The size of each kink is given in Table A1, where size is measured by the percentage change of the marginal net-of-tax rate (one minus the marginal tax rate).[2] The five largest kinks occur at gross incomes below $50,000, reflecting the strong incentives of the Earned Income Tax Credit (EITC) and Child Tax Credit (CTC). There are, however, some sizable kinks at high incomes as well. The sixth largest kink is the second statutory kink, occurring at $71,250, where statutory rates rise from 15 percent to 25 percent. In addition, there are substantial kinks at $75,000 and $113,700, at the beginning of the CTC phase-out and the threshold for Federal Insurance Contributions Act (FICA) taxes, respectively. For this taxpayer, the refund-maximizing kink is the CTC kink, where effective rates switch from negative to positive.

In addition to the four kinks of Table 1 in the main body of the paper, we also find some weak bunching at the zeroth and third kinks in the statutory schedule. The latter sees statistically significant bunching only by married taxpayers filing separately, and only in a handful of years in our sample. Panel (b) of Figure A2 depicts the full population of tax returns for this group, with data pooled from 1996 to 2014.[3] Other groups, including married couples filing jointly, fail

[1] The sizes and locations of the kinks are different for taxpayers with different filing status or tax unit size, among other factors. In particular, many of the kinks depicted do not exist for taxpayers without dependents.

[2] We measure kink size this way because the percentage change in the marginal net-of-tax rate corresponds to the denominator in the conventional definition of the elasticity of taxable income. Thus our measure of kink size has the property that, for any fixed elasticity value, larger kinks should generate larger bunching responses.

[3] At the third statutory kink and all kinks above it, we analyze the universe of tax returns because our 10 percent sample yields insufficient mass to distinguish bunching from noise.

Figure A1: Kinks faced by a single parent with two children in 2014

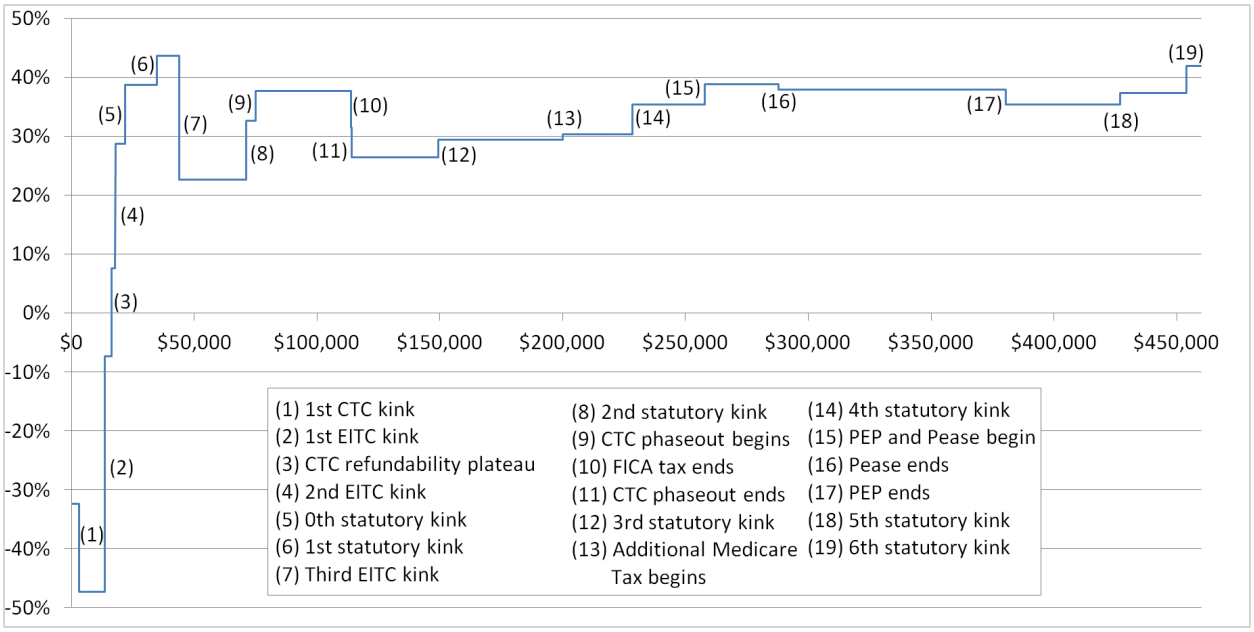

Source: authors' calculations.
The vertical axis is marginal tax rates; the horizontal axis is wage income. We assume that the taxpayer (i) only has wage income, (ii) pays no state income taxes, (iii) has $10,000 in itemized deductions, (iv) claims the EITC and CTC, and (v) does not claim the AOTC. We ignore the Alternative Minimum Tax. To measure PEP kink sizes we take the most conservative approach, assuming the marginal increment to income is $2,500. See Appendix B for a discussion of this assumption. Kinks associated with the Making Work Pay Tax Credit – applicable in 2009 and 2010 – are not pictured here. Note that the CTC kink marks the transition from negative to positive marginal tax rates and therefore is the refund-maximizing kink for this group.

to respond during any years of our sample. Curiously, among married-filing-separately taxpayers, we find stronger responsiveness among wage earners than the self-employed. The former bunch in most years of our sample, while the latter bunch in only a handful of years.

The zeroth kink presents a potential censoring issue in the data, as taxpayers without self-employment income are not required to file a tax return if their income is below this kink. In addition, taxpayers with children effectively do not face a kink here because the non-refundable portion of the CTC immediately eliminates their liability. For this reason, we only analyze childless, self-employed taxpayers at this kink. Their income distribution, centered around the kink and pooled across years 2002 to 2014, is displayed in panel (a) of Figure A2. This group produces small, statistically significant bunching in roughly half of the years in the sample.

A2. Robustness of Bunching Estimation Technique

We now test our estimation technique for sensitivity to parameter choice and functional form. The three key parameters are binwidth (δ) and the sizes of the bunching window (W) and bunching region (R). Binwidth measures how finely the data are collapsed when performing the analysis. The bunching window defines the area within which we count the total number of bunchers; we assume bunching does not occur outside the bunching window. Finally, the bunching region defines the area outside the bunching window that we use when constructing the counterfactual distribution of income if there were no kink.

Table A1: Kinks faced by a single parent with two children in 2014, ranked by size

Kink	Gross Income	Percentage Point ΔNTR	Percentage ΔNTR
Third EITC kink	$47,756	+21.1	+37.4
First EITC kink	$13,650	−40.0	−27.2
Second EITC kink	$17,830	−21.1	−22.8
Zeroth statutory kink	$21,850	−10.0	−14.0
CTC kink (maximizes refundability)	$16,333	−15.0	−14.0
Second statutory kink	$71,250	−10.0	−12.9
Beginning of CTC refundability	$3,000	+15.0	+11.3
Threshold for FICA taxes	$113,700	+06.2	+09.9
First statutory kink	$34,800	−05.0	−08.2
Beginning of CTC phase-out	$75,000	−05.0	−07.4
Sixth statutory kink	$454,050	−04.6	−07.3
End of CTC phase-out	$114,000	+05.0	+07.3
Fourth statutory kink	$228,450	−05.0	−07.2
Beginning of PEP and Pease	$257,800	−03.5	−05.5
End of PEP	$380,300	+02.5	+04.1
Third statutory kink	$149,400	−03.0	−04.1
Fifth statutory kink	$426,950	−02.0	−03.1
End of Pease	$287,800	+01.0	+01.6
Additional Medicare Tax threshold	$200,000	−00.9	−01.3

Source: authors' calculations
This table ranks the kinks of Figure A1 in descending size, measured by percentage change in the net-of-tax rate (NTR). See the caption of Figure A1 for our assumptions.

Figure A2: Bunching at two kinks

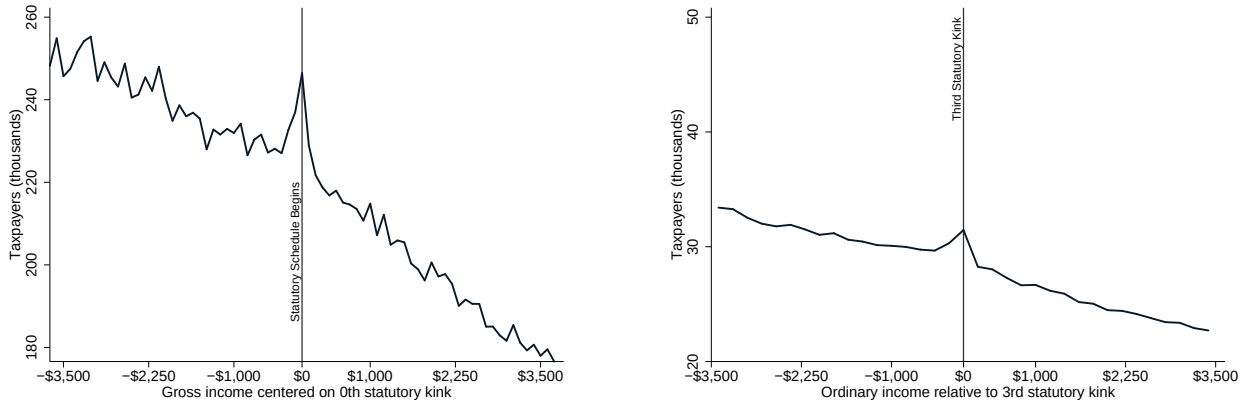

(a) 0th statutory kink: Self-employed, childless taxpayers (2002-2014)

(b) 3rd statutory kink: Married-filing-separately (1996-2014)

Panels (a) and (b) display the distribution of gross income from 2002 to 2014 for singles and married couples filing jointly, respectively, with self-employment income and no dependents. Income is rounded into $100 bins, and the distribution is centered on the zeroth statutory kink. Panel (c) displays the distribution of taxable income for the full population of married couples filing separately from 1996-2014. Income is rounded into $100 bins, and the distribution is centered on the third statutory kink. All dollar amounts are measured in 2014 dollars. All panels were created by the authors using data from the population of tax returns and have been weighted to represent the full U.S. population.

Our default parameter values are $\delta = \$150$, $W = \$1500$, and $R = \$6000$. Thus we assume that bunchers locate within \$1,500 of the kink. We use the observed distribution of income between \$1,500 and \$6,000 away from the kink to estimate the counterfactual distribution of income if the kink did not exist. We reallocate the bunching mass between \$1,500 and \$6,000 above the kink. All dollar figures here are presented in 2014 dollars.

Tables A2 and A3 test how these parameters affect our bunching coefficients for the four most responsive groups at the first EITC kink, using 2003 data. We choose this year as it is the most recent year in which self-employed, low-income taxpayers bunch *only* at the first EITC kink. Starting in 2004, we are constrained when choosing the size of the bunching region, as self-employed taxpayers bunch at the sometimes-nearby second EITC kink. We discuss this constraint in further detail in the body of the paper. By presenting 2003 estimates here, we avoid this issue and thus are able to test a wide range of parameter choices.

The results indicate that our findings are generally robust to parameter choice. For example, our preferred estimate for the bunching coefficient of the most responsive group in 2003 – single, self-employed individuals with one child – is 41.0 percent. That is, we estimate that 41.0 percent of taxpayers within the bunching window are bunchers. Binwidth choices of \$50 or \$150 lead to estimates of 41.3 percent and 40.6 percent, respectively, with small accompanying standard errors. Changing the bunching region by \$1,000 in either direction has somewhat larger effects, with alternative estimates of 36.4 percent and 44.6 percent, again with small standard errors. However, the parameter that has the most substantive effect on the estimates is the choice of bunching window. In particular, expanding the bunching window to \$2,000 causes the bunching coefficient to fall to 33.4 percent. This is unsurprising, as visual inspection of panel (a) of Figure A3 makes clear that far fewer taxpayers between \$1,500 and \$2,000 away from the kink should be classified as bunchers. Similarly, reducing the bunching window from \$1,500 to \$1,000 increases the bunching coefficient significantly, to 49.9 percent. This is also intuitive, as visual inspection clearly shows that a greater fraction of taxpayers within \$1,000 of the kink are bunchers relative to taxpayers within \$1,500 of the kink. One might have expected an even greater increase in the bunching coefficient. However, the reason the bunching coefficient does not rise more is that the counterfactual distributions are affected by the choice of bunching window. The observed distribution of income just outside of \$1,000 away from the kink appears to be elevated due to bunching, leading to artificially high estimates for the counterfactual distributions of income near the kink when a bunching window of \$1,000 is used.

Other groups show similar patterns. Our preferred bunching coefficient estimate for single, self-employed taxpayers with two children is 38.3 percent. Except for the choice of bunching window, all other permutations leave the bunching coefficient in the range $[33.8\%, 42.1\%]$. For these taxpayers a smaller bunching window of \$1,000 increases the estimate to 45.2 percent, and a larger bunching window reduces it to 31.4 percent. For married, self-employed taxpayers, our preferred bunching coefficient estimates are 28.7 percent and 28.0 percent, respectively, for those with one or two children. Except for the choice of bunching window, alternative estimates for these parameters lie in the range $[25.3\%, 30.9\%]$. For all cases, the bunching window could arguably be expanded from \$1,500 to \$2,000. This would generally decrease our bunching coefficients, but it would change the interpretation of the coefficient (from percentage within \$1,500 who are bunching to percentage within \$2,000 who are bunching) such that the drop in bunching coefficients would not substantively affect the count of total bunchers.

Tables A4 and A5 explore different functional form assumptions for the the construction of the

Table A2: Bunching coefficients calculated at the first EITC kink in 2003

	Bunching coefficient	Observations	Bin-width	Bunching window	Bunching region	Polynomial degree
Single, self-employed, one child	41.0% (1.2%)	53,800	$100	$1,500	$6,000	3
	49.9% (1.4%)	53,800	$100	$1,000	$6,000	3
	33.4% (1.2%)	53,800	$100	$2,000	$6,000	3
	36.4% (1.4%)	49,900	$100	$1,500	$5,000	3
	44.6% (1.2%)	57,100	$100	$1,500	$7,000	3
	41.3% (1.0%)	53,700	$50	$1,500	$6,000	3
	40.6% (1.5%)	53,900	$150	$1,500	$6,000	3
Single, self-employed, two children	38.3% (1.2%)	45,400	$100	$1,500	$6,000	3
	45.2% (1.6%)	45,400	$100	$1,000	$6,000	3
	31.4% (1.0%)	45,400	$100	$2,000	$6,000	3
	33.8% (1.3%)	42,600	$100	$1,500	$5,000	3
	42.1% (1.4%)	47,800	$100	$1,500	$7,000	3
	38.4% (1.0%)	45,300	$50	$1,500	$6,000	3
	38.1% (1.4%)	45,500	$150	$1,500	$6,000	3

Bunching coefficients – the percentage of taxpayers within the bunching window that are estimated to be bunching – are reported for single, self-employed, EITC-eligible filers with one or two children in 2003. Standard errors are in parentheses. The table shows the sensitivity of our estimates to variation in the estimation parameters, reported in the final four columns. Sample size reports the number of taxpayers within the bunching region and is rounded to the nearest hundred. The self-employed are those with nonzero self-employment income. Single status includes "head of household" filers. All figures are unweighted and reflect the authors' calculations using data from the population of tax returns.

Table A3: Bunching coefficients calculated at the first EITC kink in 2003

	Bunching coefficient	Observations	Bin-width	Bunching window	Bunching region	Polynomial degree
Married, self-employed, one child	28.7% (1.2%)	16,600	$100	$1,500	$6,000	3
	36.8% (1.3%)	16,600	$100	$1,000	$6,000	3
	21.8% (1.4%)	16,600	$100	$2,000	$6,000	3
	25.8% (1.3%)	14,600	$100	$1,500	$5,000	3
	30.9% (1.1%)	18,400	$100	$1,500	$7,000	3
	29.0% (1.1%)	16,600	$50	$1,500	$6,000	3
	28.5% (1.2%)	16,600	$150	$1,500	$6,000	3
Married, self-employed, two children	28.0% (1.2%)	21,000	$100	$1,500	$6,000	3
	34.4% (1.3%)	21,000	$100	$1,000	$6,000	3
	23.0% (1.1%)	21,000	$100	$2,000	$6,000	3
	25.3% (1.4%)	18,500	$100	$1,500	$5,000	3
	29.9% (1.1%)	23,300	$100	$1,500	$7,000	3
	28.1% (1.0%)	21,000	$50	$1,500	$6,000	3
	27.9% (1.3%)	21,100	$150	$1,500	$6,000	3

Bunching coefficients – the percentage of taxpayers within the bunching window that are estimated to be bunching – are reported for married, self-employed, EITC-eligible filers with one or two children in 2003. Standard errors are in parentheses. The table shows the sensitivity of our estimates to variation in the estimation parameters, reported in the final four columns. Sample size reports the number of taxpayers within the bunching region and is rounded to the nearest hundred. The self-employed are those with nonzero self-employment income. Married taxpayers who file separately are ineligible for the EITC and are excluded. All figures are unweighted and reflect the authors' calculations using data from the population of tax returns.

Table A4: Bunching coefficients calculated at the first EITC kink in 2003

	Bunching coefficient	Observations	Bin-width	Bunching window	Bunching region	Polynomial degree
Single, self-employed, one child	41.0% (1.2%)	53,800	$100	$1,500	$6,000	3
	51.4% (1.7%)	53,800	$100	$1,500	$6,000	1
	39.9% (1.7%)	53,800	$100	$1,500	$6,000	2
	35.6% (1.5%)	53,800	$100	$1,500	$6,000	4
	35.8% (1.3%)	53,800	$100	$1,500	$6,000	5
	34.2% (1.9%)	53,800	$100	$1,500	$6,000	6
	34.3% (1.7%)	53,800	$100	$1,500	$6,000	7
Single, self-employed, two children	38.3% (1.2%)	45,400	$100	$1,500	$6,000	3
	52.3% (2.1%)	45,400	$100	$1,500	$6,000	1
	37.2% (1.8%)	45,400	$100	$1,500	$6,000	2
	31.7% (1.4%)	45,400	$100	$1,500	$6,000	4
	32.0% (1.2%)	45,400	$100	$1,500	$6,000	5
	28.6% (1.6%)	45,400	$100	$1,500	$6,000	6
	28.7% (1.5%)	45,400	$100	$1,500	$6,000	7

Bunching coefficients – the percentage of taxpayers within the bunching window that are estimated to be bunching – are reported for single, self-employed, EITC-eligible filers with one or two children in 2003. Standard errors are in parentheses. The table shows the sensitivity of our estimates to variation in the estimation parameters, reported in the final four columns. Sample size reports the number of taxpayers within the bunching region and is rounded to the nearest hundred. The self-employed are those with nonzero self-employment income. Single status includes "head of household" filers. All figures are unweighted and reflect the authors' calculations using data from the population of tax returns.

Table A5: Bunching coefficients calculated at the first EITC kink in 2003

	Bunching coefficient	Observations	Bin-width	Bunching window	Bunching region	Polynomial degree
Married, self-employed, one child	28.7% (1.2%)	16,600	$100	$1,500	$6,000	3
	34.2% (1.4%)	16,600	$100	$1,500	$6,000	1
	27.9% (1.9%)	16,600	$100	$1,500	$6,000	2
	25.8% (1.9%)	16,600	$100	$1,500	$6,000	4
	26.0% (1.5%)	16,600	$100	$1,500	$6,000	5
	26.2% (2.5%)	16,600	$100	$1,500	$6,000	6
	26.2% (2.3%)	16,600	$100	$1,500	$6,000	7
Married, self-employed, two children	28.0% (1.2%)	21,000	$100	$1,500	$6,000	3
	33.4% (1.3%)	21,000	$100	$1,500	$6,000	1
	27.2% (1.6%)	21,000	$100	$1,500	$6,000	2
	25.6% (1.6%)	21,000	$100	$1,500	$6,000	4
	25.7% (1.7%)	21,000	$100	$1,500	$6,000	5
	20.4% (2.2%)	21,000	$100	$1,500	$6,000	6
	20.4% (1.9%)	21,000	$100	$1,500	$6,000	7

Bunching coefficients – the percentage of taxpayers within the bunching window that are estimated to be bunching – are reported for married, self-employed, EITC-eligible filers with one or two children in 2003. Standard errors are in parentheses. The table shows the sensitivity of our estimates to variation in the estimation parameters, reported in the final four columns. Sample size reports the number of taxpayers within the bunching region and is rounded to the nearest hundred. The self-employed are those with nonzero self-employment income. Married taxpayers who file separately are ineligible for the EITC and are excluded. All figures are unweighted and reflect the authors' calculations using data from the population of tax returns.

Figure A3: Income distribution of self-employed taxpayers near the first EITC kink in 2003

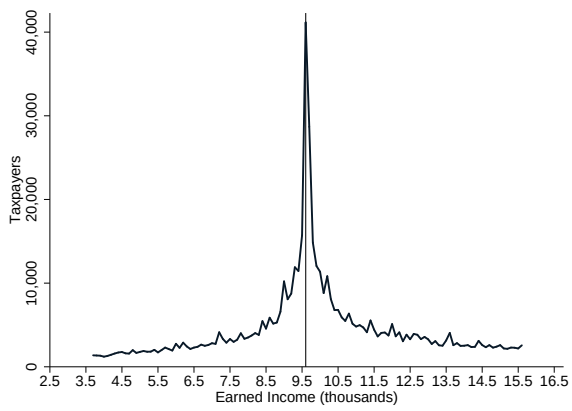

(a) Single, one child

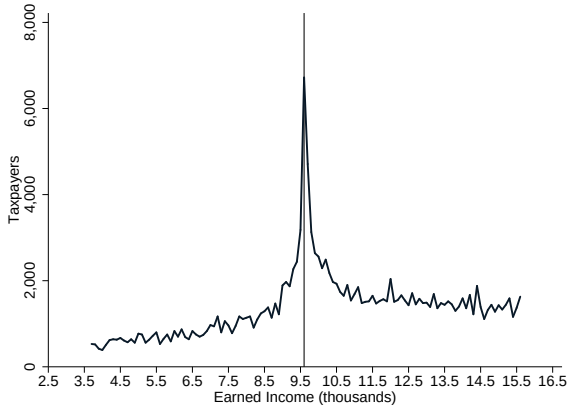

(b) Married filing jointly, one child

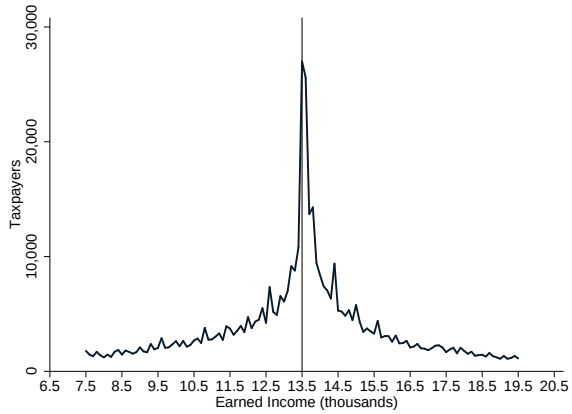

(c) Single, two children

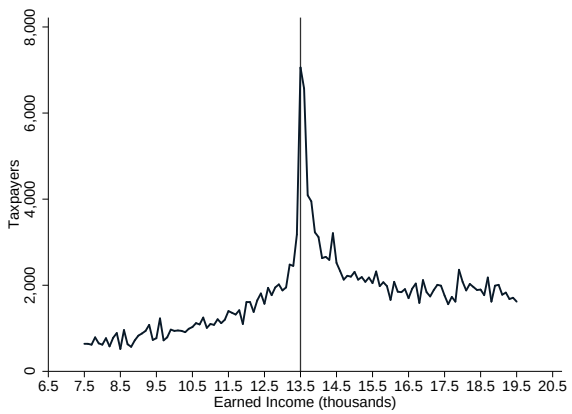

(d) Married filing jointly, two children

The distribution of income is displayed for various household types in 2003. Single status includes "head of household" filers. The self-employed are those with nonzero self-employment income. All panels were created by the authors using data from the population of tax returns and have been weighted to represent the full U.S. population. All dollar amounts are measured in 2014 dollars.

counterfactual distributions of income. Our default functional form is a polynomial of degree 3. The tables explore alternative choices for polynomial degree, ranging from one to seven. Again our default estimates are presented in the first rows of the tables. Across all four groups, a linear counterfactual density (degree 1) leads to a higher bunching coefficient, implying a greater number of bunchers than we estimate in our preferred specification. This is because the linear fit does not account for the curvature of the densities leading up to the kink, which the four panels of Figure A3 illustrate. It is not clear whether we *should* account for this curvature, as it may be a by-product of the bunching itself. However, we take the conservative approach by selecting a third-degree polynomial, resulting in a lower estimate for the number of bunchers than the first-degree polynomial would imply.

Other polynomial degrees tested (2, 4–7) generally result in smaller bunching coefficients – and thus a weaker overall magnitude of bunching – relative to the third-degree polynomial. Thus perhaps an even more conservative approach would be to use one of these polynomial degrees.

Figure A4: Alternative polynomial degrees for counterfactual estimation

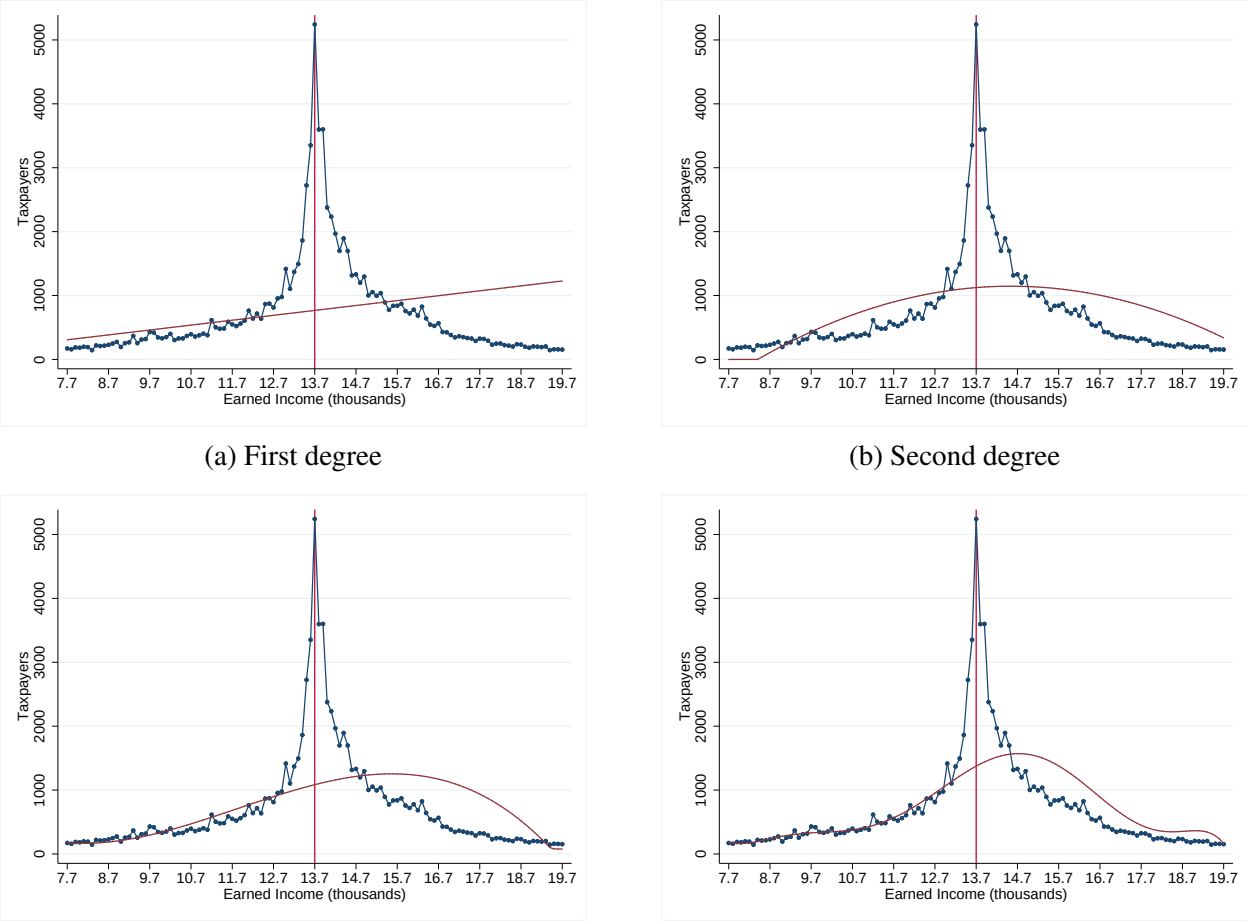

(a) First degree

(b) Second degree

(c) Third degree (preferred)

(d) Seventh degree

Actual and counterfactual distributions of income are displayed for varying polynomial degrees. The population is single, self-employed taxpayers with two children in 2014. All panels were created by the authors using data from the population of tax returns.

However, in the process of selecting our preferred specification, we observed the performance of different polynomial degrees across a wide variety of demographic groups. In our judgment, higher order polynomials tend to overfit the data. As an example, panel (d) of Figure A4 shows a seventh-degree polynomial for single, self-employed taxpayers with two children in 2014. The counterfactual density retains a relatively steep hump near the kink area. This is due to the upward curvature as the counterfactual approaches the kink, and is likely an artifact of the bunching at the kink, rather than what we would expect if the kink were removed. The third-degree polynomial featured in panel (c), in contrast, has a wider, shallower hump that we view as more realistic.

In panel (b), we see that the second-degree polynomial takes this even further, with a smooth, symmetric hump throughout the region. We view this counterfactual as implausible as well, as it features elevated mass to the left of the kink. Overall, we chose the lowest degree polynomial that we viewed as reliably producing reasonable counterfactuals across a wide variety of demographic groups and income distribution patterns.

Appendix B: Descriptive Statistics

Table B1 displays a battery of descriptive statistics for our Main Sample, with dollar values adjusted for inflation to 2014 levels. Here, as in the body of the paper, earned income refers to income that qualifies for the EITC and is equal to the sum of wages, self-employment income, and long-term disability income less one-half of self-employment "payroll" taxes. Self-employment income refers to income reported on Schedule SE. The table shows that tax units with earned income earn around $55,300 on average. This is primarily composed of wage income, which averages around $53,800 for tax units who receive it. As expected, the income distribution is highly skewed, with median incomes far below average incomes. Though the table does not report it, 12 percent of tax units in our sample have self-employment income, 18 percent receive the EITC, 56 percent use paid tax preparers, and 40 percent are married.

Table B1: Main Sample descriptive statistics (1996-2014)

	Unconditional			If greater than zero		
	Median	Mean	Std. Dev.	Median	Mean	Std. Dev.
Earned income	28,354	48,931	175,353	34,058	55,286	185,432
Wages and tips	26,031	45,005	163,159	34,395	53,804	177,066
Self-employment income	0	3,767	60,881	10,050	31,717	173,957
Long-term disability	0	349	2,704	13,640	15,086	9,683
Number of dependents	0	0.61	0.93	2	1.70	0.74
Total Observations	257,514,918					

All dollar amounts reflect taxpayer-reported figures and are inflation adjusted to 2014 levels. The number of dependents is censored such that the 6 percent of observations with three or more dependents are encoded as having three dependents. All figures reflect the authors' calculations using data from the population of tax returns.

Appendix C: Kinks That Do Not Generate Bunching

Here we describe the kinks we study that do not generate meaningful bunching patterns. These include certain state kinks as well as federal statutory kinks and effective kinks created by phase-outs of personal exemptions, itemized deductions, and the American Opportunity Tax Credit.

C.1 Third, Fourth, Fifth, and Six Statutory Kinks

Figure C1 shows the general absence of bunching at the third, fourth, fifth, and sixth kinks in the federal statutory schedule. Because the mass is so thin at these higher income levels, we use the universe of tax returns to construct this figure, rather than the ten percent sample used in the Main Sample. Panels (a), (b), and (c) pool across all years of our sample, 1996 to 2014. Panel (d) pools across 2013 and 2014, because these are the only years during our sample period when this kink exists. During our sample period, marginal tax rates changed by between two to five percentage points at these kinks. The lack of bunching is perhaps unsurprising given these relatively weak incentives.

C.2 PEP and Pease

When determining taxable income, both personal exemptions and itemized deductions phase-out at high incomes, creating discontinuities in the budget constraints of high-income taxpayers. Our evidence suggests taxpayers do not respond to these incentives, but we describe them here for completeness. The phase-outs discussed in this section were in effect during our sample from 1996 to 2005 but were gradually removed beginning in 2006, with full removal from 2010 to 2012. They have since been reinstated; however, for brevity, we do not discuss the 2013-2014 parameters.

The personal exemption phase-out (PEP) is a step function of adjusted gross income (AGI), generating notches in the budget constraint. Personal exemptions are reduced by 2 percent for each $2,500 of income exceeding the phase-out threshold until exemptions are exhausted. The beginning of the phase-out varies by filing status: $145,950 for singles, $182,450 for head of household, and $218,950 for married couples filing jointly in 2005. For all filers, the end of the phase-out region is $122,500 above the beginning.

The itemized deduction phase-out – often referred to as "Pease" after former Ohio Congressman Donald Pease – reduces certain itemized deductions at a rate of 3 cents per dollar of AGI exceeding the threshold. Pease does not apply against itemized deductions generated from casualty and theft losses, investment interest, gambling losses, or medical expenses. The total percentage of itemized deductions eliminated by Pease is capped at 80 percent per taxpayer. Throughout the time period we study (1996-2005) this threshold is the same for all filing statuses except married couples filing separately, for whom the threshold is halved. In 2005 the threshold was $145,950 ($72,975), identical to the PEP threshold for singles.

Pease creates relatively small changes in marginal tax rates at its introduction and conclusion. For example, suppose a head of household with three children claims $20,000 of itemized deductions and earns exactly the Pease threshold of $145,950 in 2005. For a marginal increase of $1,000 above the Pease threshold, qualified itemized deductions are reduced by 3 percent, meaning the individual has 30 additional dollars of taxable income. If the taxpayer faces an initial marginal tax rate of 31 percent, she would see her marginal rate increase by around 1 percentage point

Figure C1: Absence of bunching at four kinks (1996-2014)

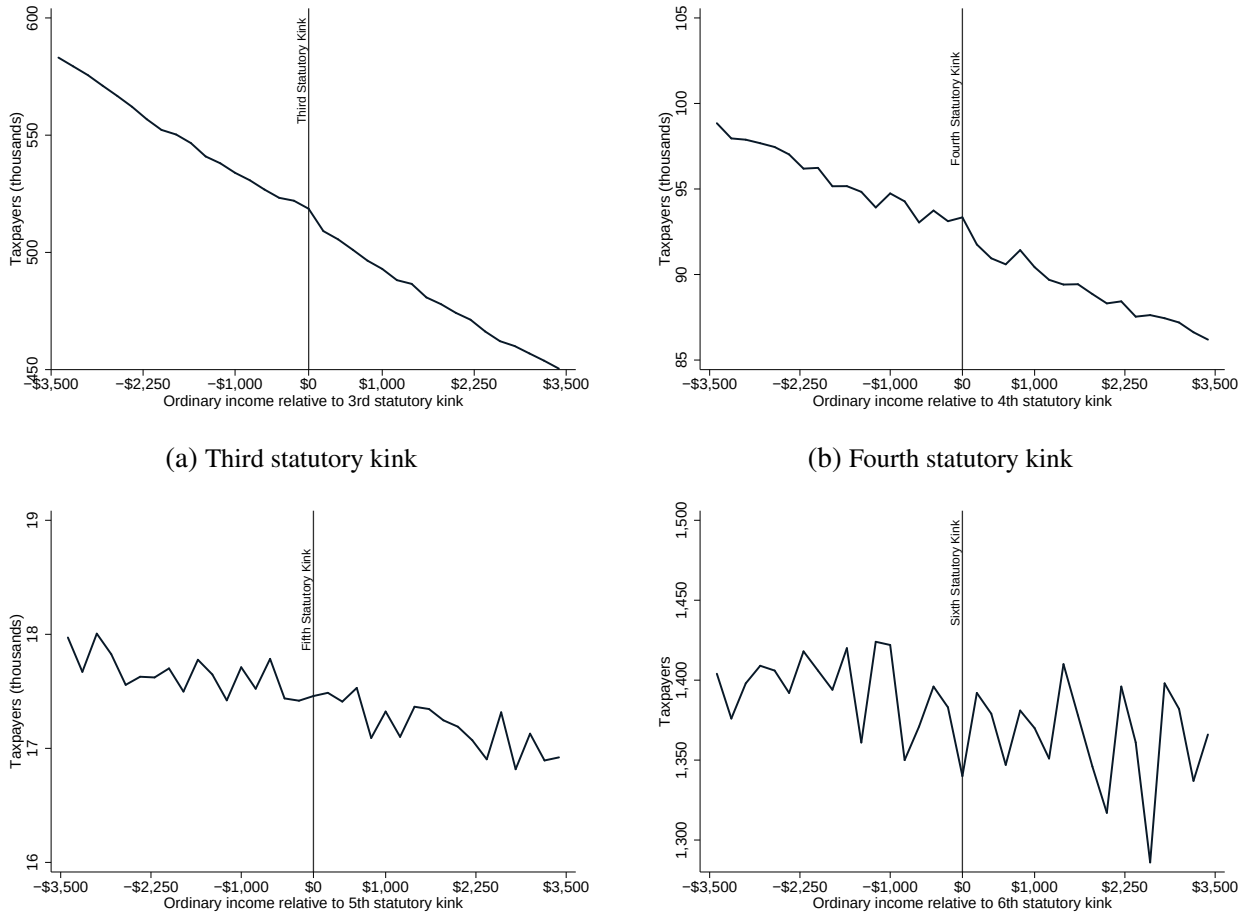

(a) Third statutory kink

(b) Fourth statutory kink

(c) Fifth statutory kink

(d) Sixth statutory kink (2013-2014)

Distributions of ordinary taxable income – centered around the third, fourth, fifth, and sixth statutory kinks – are displayed. The sixth statutory kink only existed in years 2013 and 2014 during our sample period. Bin widths are $200. All panels were created by the authors using data from the population of tax returns. Unlike in the Main Sample, here the data comprise the universe of individual tax returns and all dollars are nominal.

($\$30 \times 31\%/\$1000$) as a result of Pease, creating a small convex kink. Similarly, once Pease is phased out the change is also around 1 percentage point, which creates a small non-convex kink.

PEP generates larger marginal tax rate increases than Pease. However, because the discontinuities PEP generates are notches, not kinks, assumptions are needed to calculate the magnitude of the discontinuity relative to a kink. For example, the size of a kink is determined by the difference between the net-of-tax rates (one minus the marginal tax rate) on either side of a kink. Calculating the size of a notch requires an assumption about the incremental size of a marginal response by the taxpayer.

Suppose a taxpayer earns income at the PEP threshold. She has four personal exemptions, which reduce taxable income by $4 \times \$3,200 = \$12,800$. If this taxpayer earns at least 1 additional dollar but less than 2,500 additional dollars, her personal exemptions will be reduced by $256 ($2\% \times \$12,800$). Assuming her marginal tax rate is 31 percent initially, this increases her tax liability by $31\% \times \$256 = \$79.36$. If we assume a marginal response constitutes a $1 change in

income, the implicit change in marginal tax rates is 7,936 percentage points. If instead we assume the marginal response is $1,000, the implicit change in marginal tax rates is 7.936 percentage points. We take the most conservative measure, assuming the income increment is the full $2,500. Thus we take this taxpayer's kink size to be 3.17 percentage points.

We study bunching at the PEP and Pease kinks using the universe of tax returns from 1996 to 2005. We find no evidence of bunching for any groups in any year. Given that these kinks are fairly small, and given the lack of bunching at other high-income kinks, it is perhaps unsurprising that PEP and Pease do not generate bunching.

C.3 American Opportunity Tax Credit

The American Opportunity Tax Credit (AOTC) is a partially refundable tax credit for qualified post-secondary education expenses. The maximum credit amount is $2,500 per eligible student, which is achieved by having at least $4,000 in qualified expenses per eligible student. The credit is refundable at a 40 percent rate, up to a maximum of $1,000 per eligible student. The AOTC creates kinks at two points: a convex kink at the beginning, and a non-convex kink at the end of the phase-out region. For single-headed tax units the phase-out begins at $80,000 of modified adjusted gross income (MAGI), and ends at $90,000. For married-filing-jointly taxpayers, the phase-out begins at $160,000 and ends at $180,000. These phase-out regions create effective marginal tax rate vary by number of students with qualifying expenses. A single-headed tax unit with two student dependents who each have qualifying expenses exceeding $4,000 experiences an effective marginal tax rate increase of 50 percentage points at $80,000 of MAGI. The analogous marginal tax rate increase for married tax units is 25 percentage points at $160,000, as the phase-out region is twice as long (ending at $180,000).

We study bunching at AOTC kinks using the universe of tax returns from 2009 to 2014 that claim the AOTC. We study only the convex kinks marking the beginning of the phase-out, as taxpayers just above the end of the phase-out do not claim the credit and therefore we cannot cleanly identify the relevant population of eligible taxpayers. For most taxpayers, in most years, there is no bunching response at the convex kink. For the few tax units types that do appear to bunch in certain years, the magnitude of the response is economically trivial. Given the size of these kinks, it is surprising they do not provoke a meaningful bunching response. In many cases, the convex AOTC kink causes increases in marginal tax rates that are similar in magnitude to the EITC and CTC kinks. Thus it provides an opportunity to test the hypothesis that high-income taxpayers would bunch if they faced incentives similar to the EITC or CTC. However, an important difference between the AOTC and the EITC or CTC is that the AOTC does not mark the refund-maximizing kink for any taxpayers in any year.

C.4 State Kinks

In a further search for responsiveness to tax kinks among high-income taxpayers, we tested for bunching at a few kinks marking the top tax bracket in state tax schedules. California had a kink that increased marginal tax rates by one percentage point at $1 million of taxable income during 2006 to 2014. For most years, this kink applied regardless of filing status; however, in 2011 the kink was $2 million for married couples filing jointly. Connecticut had a kink that increased marginal tax rates by 1.5 percentage points at $500,000 ($1 million for married couples filing

jointly) during 2010 and 2011. Finally, New Jersey had a kink that increased marginal tax rates by between 1.6 to 2.5 percentage points at $500,000 during 2004 to 2014 ($1 million for married couples filing jointly during 2004 to 2010).

While these kinks are fairly small compared to federal statutory kinks, they offer a change to test for bunching at incomes far exceeding any kinks in the federal schedule. We study the state kinks listed above using the universe of tax returns filed with the IRS by taxpayers in these states during the years in which the kinks were active. We see no evidence of bunching among any groups in any years at these kinks.[4]

[4]One caveat is that New Jersey's definition of taxable income is different than the federal government's. Thus it remains possible that taxpayers bunched at New Jersey's high-income kink and we are unable to detect it.