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

We document bunching at tax kinks using a panel of 258 million income tax returns in the United States from 1996 to 2014. During this period bunching grew by 700%. While most bunchers are self-employed, a substantial number of wage earners also bunch by misreporting income. The vast majority of bunching occurs at kinks maximizing tax credits, particularly at the kink that maximizes taxpayer refunds. Many taxpayers follow the refund-maximizing kink from year to year, selectively maximizing the appropriate tax credit(s). We argue that this behavior is incompatible with recently developed methods for identifying elasticities via bunching patterns.
(*JEL* H23, H24, H26, H31, J22)

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

*For their helpful comments, we thank Gerald Auten, Thomas Barthold, Claire Brunel, James Cilke, Timothy Dowd, Nada Eissa, John Friedman, Robert Harvey, Patrick Langetieg, Arik Levinson, Robert McClelland, John Rust, and Heidi Schramm; seminar participants at the Congressional Budget Office, Georgetown University, Institute for Social Research, Pennsylvania State University–Harrisburg, Tax Economists Forum, Tax Policy Center, University of South Dakota, U.S. Department of the Treasury, and the U.S. Department of Justice; and attendees of the meetings of the American Economic Association, Association for Public Economic Theory, Georgetown Center for Economic Research, National Tax Association, Society for Government Economists, and the Zentrum für Europäische Wirtschaftsforschung. Andrew gratefully acknowledges the Economic Club of Washington, D.C. for financial support. Previous versions of this paper circulated under the title "How Sensitive Are Taxpayers to Marginal Tax Rates? Evidence from Income Bunching in the United States." The views expressed in this paper are our own and are not necessarily those of the U.S. Department of the Treasury. This research embodies work undertaken for the staff of the Joint Committee on Taxation, but as members of both parties and both houses of Congress comprise the Joint Committee on Taxation, this work should not be construed to represent the position of any member of the Committee. The authors declare that they have no relevant financial interests that relate to the research described in this paper.

This paper studies taxpayer bunching at kinks in tax schedules. A kink is an income amount for a given taxpayer at which marginal tax rates change discretely, marking the end of one tax bracket and the beginning of the next. Standard economic theory predicts that some taxpayers will avoid brackets with relatively high tax rates by bunching at kinks where tax rates increase, resulting in extra mass in the distribution of income close to these kinks. We measure this excess mass at a wide variety of kinks in the United States federal tax schedule, building upon methods developed by Saez (2010) and Chetty et al. (2011).

Understanding the way taxpayers respond to the tax code is necessary for welfare and revenue analyses of current and proposed income tax regimes. There are two main ingredients for such policy analysis: a model of behavior and data to calibrate the model. As Einav, Finkelstein and Schrimpf (2016) illustrate, given the same data, different model choices can lead to different forecasts of behavior. The present paper sheds light on the correct model(s) behind taxpayer bunching, uncovering patterns that are incompatible with standard labor supply models of behavior.

Economists have long recognized that, under a range of assumptions, kinked budget sets should induce bunching in income distributions. However, the first searches for bunching relied upon survey data and failed to find responses at kinks. Thus, the focus of early work estimating responsiveness to the budget set was on fitting models that could identify behavioral responses to marginal tax rates but allowed for a lack of bunching at kinks. Initially this involved parametric assumptions regarding the error term in individuals' labor supply (Burtless and Hausman, 1978); however, more recent work allows for nonparametric identification (Blomquist and Newey, 2002). These nonlinear budget set models can be thought of as using responses away from kinks to identify tax rate sensitivity, while using bunching (or the lack thereof) to calibrate optimization frictions.

In contrast, Saez (2010) develops a technique for using bunching to directly estimate responsiveness to marginal tax rates at a given kink. The bunching approach relies on estimating counterfactual income distributions that would hold if there were no kink. Saez uses this method to estimate elasticities of taxable income, analyzing kinks in the Earned Income Tax Credit (EITC) and statutory federal income tax schedules using public-use tax data from 1960 to 2004. He finds substantial bunching only around the first EITC kink and at $0 of taxable income. All other kinks appear to generate no bunching. Importantly, when looking closer at the bunching patterns around the first EITC kink, Saez finds the response is driven entirely by the self-employed. Building on this work, Chetty et al. (2011) offer an alternative approach to constructing counterfactual income distributions. Our technique borrows aspects from both of these approaches.

In the United States, the bunching approach has also been used to analyze responsiveness to the Annual Earnings Test for Social Security income (Burtless and Moffitt, 1984; Friedberg, 2000; Gelber, Jones and Sacks, 2013) and the Saver's Credit notch (Ramnath, 2013). Elsewhere, bunching has been studied at discontinuities in tax schedules in Denmark (Le Maire and Schjerning, 2013), Sweden (Bastani and Selin, 2014), Pakistan (Kleven and Waseem, 2013), Ireland (Hargarden, 2015), and the United Kingdom (Devereux, Liu and Loretz, 2014). Kleven (2016) discusses this research as well as several applications of the bunching approach outside of the tax literature.

Our bunching measures are derived from nationally-representative panel data, drawn from the universe of federal income tax returns in the United States from 1996 to 2014. With 258 million observations in total, most of our estimators – including those for narrowly defined household types in a given year – use tens or hundreds of thousands of observations, resulting in smooth distributions over the intervals surrounding kinks. These data cover two orders of magnitude more taxpayers than the public-use files, allowing us to study heterogeneity in responsiveness across household types and across kinks. We study all statutory kinks in the federal tax schedule as well as several effective kinks created by the EITC, Child Tax Credit (CTC), Making Work Pay Tax Credit, American Opportunity Tax Credit, and phase-outs of itemized deductions and personal exemptions.[1]

We make five contributions to a better understanding of bunching behavior. First, we show that many taxpayers gravitate towards the unique point of the tax schedule that maximizes the refund they receive. We argue that this behavior is incompatible with recently developed methods that use bunching patterns to identify the elasticity of taxable income. Second, we document the emergence of bunching by wage earners and show that this response is driven exclusively by income misreporting. Third, we show that bunching is persistent: taxpayers near kinks are disproportionately likely to have similar incomes the following year. Fourth, we measure growth in bunching – including the emergence of responsiveness at new kinks – over a nineteen-year window that includes two tax reforms and the Great Recession. Finally, we derive a new bunching estimator that combines the flexibility of the Chetty et al. (2011) technique and the applicability to large, discrete kinks of the Saez (2010) technique.

Though we test for bunching at several middle- and high-income kinks, we find strong bunching patterns only at low-income kinks. In particular, taxpayers gravitate towards kinks that maximize refundable tax credits, including the EITC, CTC, and Making Work Pay Tax Credit, suggesting taxpayers may be more attuned to incentives regarding transfer payments than to incentives regarding taxes owed. An alternative explanation is that taxpayers respond to these kinks simply because they represent the largest kinks in the budget set. The kinks created by the EITC typically change marginal tax rates by more than twenty percentage points, and the CTC kink where we see responsiveness changes rates by fifteen percentage points. However, the Making Work Pay Tax Credit changed rates by only about six percentage points – significantly less than three kinks in the statutory schedule – and it nonetheless managed to attract a significant portion of bunching during the years it was available. In other cases, too, bunching responses are oddly disproportionate to the size of the kink. We hypothesize that this is due to a portion of taxpayers gravitating towards the point in the schedule that maximizes the refund they receive.

Married, self-employed taxpayers with two children provide clear evidence of this phenomenon. In 2010, the Making Work Pay Tax Credit and the first EITC kink were nearby – within $1,000

[1] We also examine notches related to Medicaid and federal disability insurance, as well as a few high-income kinks in state tax schedules and a kink related to the Supplemental Nutritional Assistance Program. We find no response at any of these discontinuities in marginal incentives. See Appendices C and D for further details.

of each other. The former changed tax rates by six percentage points; the latter by forty percentage points. Nonetheless, bunching was stronger at the former kink, which earned taxpayers the maximum possible refund (given fixed withholding decisions). In 2011, two things changed: the Making Work Pay Tax Credit no longer existed, and payroll taxes were temporarily reduced by two percentage points. These relatively minor changes moved the location of the refund-maximizing kink to the CTC refundability plateau, which begins the region where the CTC is maximized. Accordingly, the distribution of income for this group saw a major shift, such that bunching now peaked at the CTC kink. The distribution changed little in 2012, but in 2013 the payroll tax holiday expired and the refund-maximizing kink changed to the first EITC kink. A large portion of taxpayers responded immediately, with bunching plummeting at the CTC kink and rising dramatically at the first EITC kink.[2]

The behavior we have just described – with taxpayers targeting maximum refunds – is incompatible with the models of taxpayer behavior used to translate bunching patterns into elasticities of taxable income by Saez (2010) and Chetty et al. (2011), among others. Their elasticity estimates hinge upon the assumption that the number of bunchers at a kink is proportional to the kink's size, all else equal. When taxpayers instead seek out the refund-maximizing point in the tax code, this assumption fails. These refund-maximizing taxpayers are not responsive to the percentage change in their marginal tax rates, as the concept of an elasticity presupposes. Instead, they exclusively respond to the location in the tax schedule where effective marginal tax rates switch from negative to positive. Our results should give pause to those estimating elasticities using observed bunching intensity in any context where agents can alter the variable of interest through misreporting.

At refund-maximizing kinks and elsewhere, we find that taxpayer bunching is growing rapidly. In 1996, we estimate that approximately 104,000 taxpayers bunched at the kinks we study.[3] Bunching was concentrated almost exclusively at the first EITC kink, which marked the maximum possible refund for most taxpayers. Over the next two decades, bunching grew steadily at the first EITC kink and spread to other locations in the tax schedule, most notably at the kink that maximizes the CTC. By 2014, we estimate there were 834,000 bunchers. This represents a 700% increase in bunching during a period when the tax filing population grew by only 20%.

Taxpayers' increasing tendency to report incomes at the refund-maximizing kink is especially remarkable given that relatively sophisticated calculations are needed to determine the necessary income level. Two prime candidates for informing taxpayers of bunching incentives are paid preparers and tax-filing software. However, probit regression analysis indicates that these covariates are only weakly correlated with the decision to locate near a given kink. An alternative hypothesis, which receives strong support in our data and in that of Chetty, Friedman and Saez (2013), is that taxpayers learn about bunching incentives from their neighbors. As we and Chetty, Friedman and Saez (2013) show, there is substantial geographic heterogeneity in bunching, and bunching appears

[2]The patterns described in this paragraph are on display in Figure 8.

[3]This is a subset of the total mass at these kinks, as it does not include taxpayers who would remain at these income levels if there were no kinks.

to spread across the nation slowly, emanating from the U.S. South.[4]

Regardless of how taxpayers learn about bunching incentives, the dominant mechanism taxpayers use to bunch is by reporting self-employment income. Though the self-employed represent only 12% of taxpayers overall, we estimate they comprise 89% of bunchers. This raises the question of whether bunching reflects distortions of real economic behavior or merely reporting behavior, as the self-employed are known to exhibit high rates of noncompliance (Slemrod, 2007; Internal Revenue Service, 2016). We cannot definitively answer this question; however, we do find that self-employed taxpayers who lack third-party reports of earnings are disproportionately likely to locate near kinks. Even more starkly, among wage earners, bunching patterns manifest *only* in taxpayer-reported earnings. When we plot income distributions using employer-reported income (supplemented with taxpayers' statements of unreported wage and tip income), bunching disappears. Thus, wage-earning bunching is unmistakably the product of misreported income. We caution, though, that our data are pre-audit: we observe only what taxpayers report at the time of filing. We cannot say whether the tax credits claimed are ever paid out, nor whether they are recovered by the Internal Revenue Service in the event they were paid out but later deemed inaccurate.

I Data and Institutional Background

In this section we describe our data as well as the portions of the U.S. tax code where we find bunching responses. We begin with an overview of our data sources, followed by a discussion of summary statistics. We conclude by describing the statutory federal schedule as well as the Earned Income Tax Credit, Child Tax Credit, and Making Work Pay Tax Credit.[5]

I.A Data

Our analysis of taxpayer bunching uses data drawn from the Internal Revenue Service's Compliance Data Warehouse (CDW). The CDW contains the universe of tax returns (filed by taxpayers) and information returns (typically filed by third parties) of individuals in the United States. Each observation in our data is a tax unit – a rough proxy for a household – that filed a tax return for a given year. The tax return data give various sources of income as well as deductions, credits, taxes paid, and demographic information such as marital status, number of children, address of residence, and years of birth of those in the tax unit. All data derived from tax returns are pre-audit and therefore reflect what taxpayers report when filing, including any errors. From information returns, we observe employer-reported wage income (Form W-2), independent contractor income (Form 1099-MISC), and long-term disability payments (Form SSA-1099). Finally, we also use

[4]See Appendix E for graphical evidence of the spread of bunching.

[5]In Appendix D, we describe the kinks we study that do not generate bunching, including a few high-income state kinks as well as federal kinks created by phase-outs of personal exemptions, itemized deductions, and the American Opportunity Tax Credit.

information on date of birth and sex at the time of birth from the Social Security Administration's Data Master File.

For most of the analysis, we use a panel of taxpayers that is representative of the tax-filing population in the United States in every year from 1996 to 2014. We refer to this sample as the Main Sample. Each year of data in the Main Sample contains a random selection of tax units with a 10% sampling probability based on a unique, time-invariant identifier for the primary filer (i.e., the taxpayer listed first on the return). The panel is unbalanced, as taxpayers may be primary filers one year but not the next. This causes complications when we study whether the same tax units persist at kinks year after year. For this reason, when we track tax units over time, we supplement the Main Sample with an auxiliary dataset that includes all observations of tax units whose secondary filer is a primary filer in at least one year in our Main Sample. We refer to this augmented sample as the Combined Sample.

For certain high-income kinks, where the distribution of income is thin, 10% of the tax-filing population is insufficient to distinguish bunching patterns from noise. For these kinks we use the full population of tax units in the neighborhood of the kink. These include the four highest-income kinks in the federal statutory schedule, and effective kinks created by phaseouts of personal exemptions and itemized deductions. We also use the full population of tax returns from 2003 to 2014 in California, Connecticut, and New Jersey in the neighborhood of each state's largest kink.

Table 1 displays a battery of descriptive statistics for the EITC-eligible and non-EITC-eligible observations in our Main Sample. As discussed in Section I.C, EITC eligibility is based on age, earned income, unearned income, and number of children. Earned income is comprised of wages, self-employment income, and long-term disability income. The table shows that EITC eligibility is associated with having less earned income and more children, consistent with the credit's structure, which targets low-income working families with children. The average and median EITC-eligible tax units are headed by a younger individual that is more likely to be single and female, is more likely to use a paid preparer, and is more likely to have self-employment income than non-EITC-eligible tax units.[6]

I.B Federal Tax Code

Despite its reputation for complexity, the U.S. federal income tax code has a straightforward statutory schedule. In 1996, the first year in our sample, the schedule for ordinary income had five tax brackets whose marginal tax rates are detailed in Table 2.[7] This schedule remained stable on an inflation-adjusted basis until the Bush Tax Cuts of 2001-2003, which added a 10% bracket at the

[6]Throughout the paper, we take the term "single" to refer to unmarried taxpayers, including those who elect the "head of household" filing status.

[7]The actual implementation of Table 2's marginal tax rates involves a large number of $50 micro-brackets, with discrete changes in tax liability only at the beginning of each bracket. Hence, the effective tax rate on marginal income is actually zero for most taxpayers for small enough marginal income increments. Like most other researchers, we ignore this nuance, assuming taxpayers' marginal decisions involve large dollar increments. For evidence of bunching within micro-brackets, see Slemrod (1985).

Table 1: Main Sample descriptive statistics (1996-2014)

	Unconditional		Greater than zero	
Non-EITC-Eligible	Median	Mean	Median	Mean
Earned income	42,960	63,555	48,968	70,810
Wages and tips	39,930	58,659	48,031	68,014
Self-employment income	—	4,821	10,731	41,351
Fraction self-employed	—	0.11	—	—
Wages & tips not reported by employer	—	6	1,222	3,756
Long-term disability	—	304	15,575	16,843
Age of primary	45	45	—	—
Fraction using paid preparer	—	0.55	—	—
Fraction male	—	0.73	—	—
Fraction married (filing jointly)	—	0.48	—	—
Number of dependents	0	0.48	2	1.79
Non-EITC-eligible observations		174,263,600		

	Unconditional		Greater than zero	
EITC-eligible	Median	Mean	Median	Mean
Earned income	18,391	23,540	19,458	24,616
Wages and tips	15,742	21,049	19,781	24,576
Self-employment income	—	2,044	9,839	12,714
Fraction self-employed	—	0.16	—	—
Wages & tips not reported by employer	—	10	1,421	3,285
Long-term disability	—	590	12,744	14,358
Age of primary	38	39	—	—
Fraction using paid preparer	—	0.61	—	—
Fraction male	—	0.49	—	—
Fraction married (filing jointly)	—	0.21	—	—
Number of dependents	1	1.20	1	1.61
EITC-eligible observations		59,848,895		

All dollar amounts reflect taxpayer-reported figures and are inflation adjusted to 2014 levels. EITC eligibility is calculated annually. The number of dependents is censored at three. All observations where the primary filer's age is less than twenty or greater than eighty are dropped. All figures reflect the authors' calculations using data from the IRS Compliance Data Warehouse.

beginning of the schedule and generally lowered rates. The Bush tax rates remained in place until the American Taxpayer Relief Act of 2012, which reinstated a top bracket of 39.6%.

The kink points separating the federal tax brackets vary by year and filing status. To keep terminology uniform, throughout the paper we take the "first" kink to be the divider between the

Table 2: Federal ordinary income tax: Statutory marginal tax rates (%)

	Bracket						
Year(s)	1st	2nd	3rd	4th	5th	6th	7th
1996-2000	—	15	28	31	36	39.6	—
2001	—	15	27.5	30.5	35.5	39.1	—
2002	10	15	27	30	35	38.6	—
2003-2012	10	15	25	28	33	35	—
2013-2014	10	15	25	28	33	35	39.6

Source: Section 1 of the U.S. Internal Revenue Code.
Locations of the kinks are adjusted for inflation annually. See Figure 1 for a graphical depiction of the schedule in 2014, including effective kinks not described in this table.

first and second brackets according to the post-2001 schedule. Similarly, we take the "second" kink to be the divider between the second and third brackets, and so on. Thus, in our terminology, the first kink did not exist in 1996-2001 and the sixth kink did not exist in our sample until 2013. The "zeroth" kink marks the beginning of the schedule in all years.

Unlike the EITC and CTC schedules detailed below, the statutory income tax schedule is progressive. All kinks see marginal tax rates increase and are therefore convex. Most of these kinks are small. In the presence of significant optimization frictions, we might not expect bunching at such kinks. Two of the statutory kinks, however, create absolute changes in tax rates of ten percentage points or more: the zeroth and second statutory kinks. All else equal, we expect to observe stronger responsiveness at these kinks. Estimating bunching at the zeroth kink requires care, however, as it also represents the filing threshold for most taxpayers. That is, most individuals with taxable income below this kink are not required to file federal income tax returns, potentially creating a censoring problem in the data. We avoid this issue by only examining self-employed taxpayers, as their filing threshold is $400 during our sample period. We further limit this sample to those taxpayers with no dependents to abstract away from the non-refundable portion of the CTC, which effectively eliminates the zeroth kink for most taxpayers with children.

I.C Earned Income Tax Credit

The EITC is one of the largest poverty alleviation policies (and tax expenditures) in the United States, with some 28.8 million low-income tax units receiving $68 billion in 2013.[8] These figures have grown since 1996, when roughly 19.5 million tax units received $28.8 billion.[9] All low-income taxpayers between the ages of 25 and 64 are eligible, and the age restriction applies only to taxpayers with no qualifying children. The credit's schedule varies based on filing status (single

[8] See Eissa and Hoynes (2011) and Nichols and Rothstein (2015) for detailed discussions of the EITC.

[9] These figures are taken from the Statistic of Income's "Tax Stats" website, specifically the section on the EITC here: http://www.irs.gov/Individuals/Earned-Income-Tax-Credit-Statistics.

or married), number of qualifying children, and tax year. Childless households may qualify for a small credit (a maximum of $496 in 2014), but the EITC is substantially more generous for households with children. For example, a taxpayer with three qualifying children in 2014 could potentially receive a credit of $6,143.

The term "earned" in the credit's title refers to the definition of income to which the credit applies: wage and self-employment income (as well as long-term disability payments). As earned income increases from zero, all households face a phase-in region, a plateau, and a phase-out region. In the phase-in region, earned income is subsidized at rates between 7.65% and 45% depending on the number of qualifying children in the household. In the plateau region, the taxpayer receives the maximum credit amount and the EITC does not affect marginal incentives. In the phase-out region, the subsidy is removed at rates between 7.65% and 21.06%, again depending on the number of qualifying children. This phase-out effectively increases tax rates in this region, potentially discouraging earnings.[10]

We expect individuals to bunch around the first and second EITC kinks, marking the beginning and end of the plateau region. Because it is non-convex, the kink at the end of the phase-out region should induce an absence of mass. Previous studies have not identified responses at non-convex kinks, but for many taxpayers this kink creates the largest percentage change in their net-of-tax rate (one minus the marginal tax rate). Thus, if we see a response to any non-convex kink, we expect it here.

I.D Child Tax Credit

The Child Tax Credit (CTC) is available to taxpayers on a per child basis, but phases out at high incomes.[11] The credit has both refundable and non-refundable components, with the refundable portion (called the Additional CTC) phasing in as a function of earned income. The credit amount and refundability parameters have varied since the credit's introduction in 1997. Initially the CTC was $400 per qualifying child. In 1999 it increased to $500; in 2001 and 2002 the credit was $600; and in 2003 the credit increased to its present value of $1,000 per qualifying child.

The credit creates two convex and two non-convex kinks. The first non-convex kink is the refundability threshold, which was introduced in 2001. Above this threshold the portion of the credit exceeding the taxpayer's liability can be claimed by the individual, but only at a rate of 10% or 15% of earned income exceeding the threshold, depending on the year. This has no effect on households whose tax liability exceeds the credit amount. For most households, however, the kink effectively decreases marginal tax rates by ten or fifteen percentage points. The threshold was $10,000 from 2001 to 2007 (indexed to inflation beginning in 2002), reduced in 2008 to $8,500, and reduced again in 2009 to $3,000 (no longer indexed to inflation). The refundability rate was

[10] Phase-out of the EITC occurs using the greater of earned income and adjusted gross income (AGI). We ignore this in our empirical analysis, assuming all taxpayers have weakly greater earned income than AGI. Our measures, therefore, likely understate responsiveness at the second EITC kink.

[11] See Crandall-Hollick (2013) for a full description of the CTC, including its legislative history.

10% from 2001 to 2003 and 15% after.

The first convex kink occurs at the point where the CTC becomes fully refundable. After this point the credit is maximized, creating a plateau region. The kink at the beginning of the refundability plateau is comparable to the first EITC kink, where the credit is maximized and creates a plateau. Given that taxpayers are known to bunch at the first EITC kink, the CTC refundability plateau kink is perhaps the most likely place to find a bunching response to the CTC.

The second convex kink marks the end of the credit's plateau and the beginning of the phase-out region, which is $75,000 for singles and $110,000 for married filing jointly (neither is indexed to inflation). The credit is reduced by $50 for every additional $1,000 in modified adjusted gross income (MAGI), effectively increasing marginal tax rates by five percentage points. The end of the phase-out region, where the credit is completely eliminated, marks the second non-convex kink. At this point the taxpayer's marginal tax rate decreases by five percentage points. In a frictionless world, we would expect an absence of mass at this point. However, given the size of the kinks, we do not expect to find responsiveness at either the beginning or end of the CTC phaseout region.

I.E Making Work Pay Tax Credit

The Making Work Pay Tax Credit (MWPTC) was a refundable tax credit available to low- and middle-income workers in 2009 and 2010. The credit was administered through a reduction in withholdings on Form W-2, and as a result many low-income individuals received the credit even without filing a tax return. The credit effectively reduced the tax rate on earned income by 6.2% up to $6,451 of earned income for singles and $12,903 of earned income for married couples filing jointly. The maximum credit amounts were $400 and $800 for singles and married couples, respectively. The credit began to phase-out at a rate of roughly 2% at $150,000 of MAGI for married couples filing jointly and $75,000 for all others, and was fully exhausted at $190,000 and $95,000, respectively.

The end of the phase-in region, beginning of the phase-out region, and end of the phase-out region all created kinks. However, we only expect responsiveness at the first kink, as it was a relatively large, convex kink: 6.2 percentage points for all returns. Further, the first MWPTC kink maximized a refundable tax credit and was a salient component of the American Recovery and Reinvestment Act of 2009. In contrast, the other MWPTC kinks were relatively small, creating incentives that would be easily overwhelmed by moderate optimization frictions.

II Bunching Analysis

We now turn to documenting bunching patterns at the kinks described in the previous sections. Figure 1 depicts the kinks we study for an unmarried taxpayer with two children in 2014.[12] The

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

Figure 1: Kinks faced by a single parent with two children in 2014

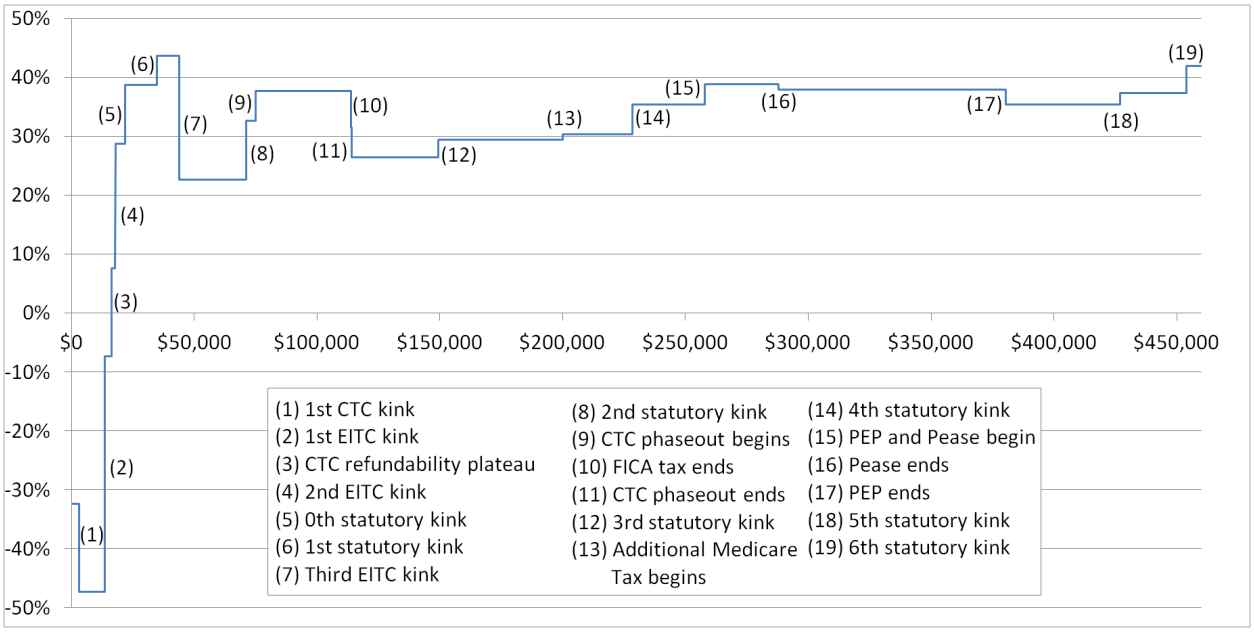

Source: authors' calculations.
The vertical axis is marginal tax rates; the horizontal axis is wage income. We assume that the taxpayer (i) only has wage income, (ii) pays no state income taxes, (iii) has $10,000 in itemized deductions, (iv) claims the EITC and CTC, and (v) does not claim the AOTC. We 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 D for a discussion of this assumption. FICA refers to the Federal Insurance Contributions Act. Kinks associated with the Making Work Pay Tax Credit – applicable in 2009 and 2010 – are not pictured here. Note that the CTC refundability plateau kink marks the transition from negative to positive marginal tax rates and therefore is the refund-maximizing kink for this group.

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 3, where size is measured by the percentage change of the marginal net-of-tax rate (one minus the marginal tax rate).[13] The five largest kinks occur at gross incomes below $50,000, reflecting the strong incentives of the EITC and 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% to 25%. In addition, there are substantial kinks at $75,000 and $113,700, at the beginning of the CTC phase-out and the threshold for FICA taxes, respectively. For this taxpayer, the refund-maximizing kink is the CTC refundability plateau, where effective rates switch from negative to positive.

If taxpayers are sensitive to marginal tax rates and optimization frictions are small, all convex kinks will generate bunching (Saez, 2010). However, at most kinks there is no evidence of responsiveness in any of the years of our sample. This includes the largest kink in Figure 1, the non-convex kink at the end of the EITC phase-out region (the seventh kink in the figure). Standard theory predicts a dip in the distribution of income near a non-convex kink. We do not observe this

[13] 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.

Table 3: Kinks faced by a single parent with two children in 2014, ranked by size

Kink	Gross Income	Percentage Point ΔNTR	Percentage ΔNTR	Response for some group in our sample?
Third EITC kink	$47,756	+21.06	+37.41	×
First EITC kink	$13,650	−40.00	−27.15	✓
Second EITC kink	$17,830	−21.06	−22.80	✓
Zeroth statutory kink	$21,850	−10.00	−14.03	✓
CTC refundability plateau	$16,333	−15.00	−13.97	✓
Second statutory kink	$71,250	−10.00	−12.93	✓
Beginning of CTC refundability	$3,000	+15.00	+11.33	×
Threshold for FICA taxes	$113,700	+06.20	+09.94	×
First statutory kink	$34,800	−05.00	−08.16	×
Beginning of CTC phase-out	$75,000	−05.00	−07.42	×
Sixth statutory kink	$454,050	−04.60	−07.34	×
End of CTC phase-out	$114,000	+05.00	+07.29	×
Fourth statutory kink	$228,450	−05.00	−07.18	×
Beginning of PEP and Pease	$257,800	−03.52	−05.45	×
End of PEP	$380,300	+02.53	+04.08	×
Third statutory kink	$149,400	−03.00	−04.08	✓
Fifth statutory kink	$426,950	−02.00	−03.09	×
End of Pease	$287,800	+00.99	+01.62	×
Additional Medicare Tax threshold	$200,000	−00.90	−01.28	×

Source: authors' calculations
This table ranks the kinks of Figure 1 in descending size, measured by percentage change in the net-of-tax rate (NTR). See the caption of Figure 1 for our assumptions. Note that the third statutory kink sees bunching only among married taxpayers filing separately, for whom the kink lies below $75,000 in all years.

in the data, even for groups that are highly sensitive to other kinks.

We also see no response at most statutory kinks, including all high-income kinks. Most surprising was the lack of meaningful bunching at the American Opportunity Tax Credit kinks, which are quite large (comparable to, and sometimes larger than, the EITC kinks) and were present from 2009 to 2014.[14] In addition to exploring federal high-income kinks, we analyze the largest high-income kinks created by state tax regimes, which occur in California, Connecticut, and New Jersey. These kinks are small relative to federal kinks (less than three percentage point changes), but occur at incomes greatly exceeding any federal kink (up to $2 million). We find no evidence of bunching at any of these kinks during any of the years in our sample period. All of the bunching patterns we observe occur at incomes below $75,000, and the strongest patterns occur at kinks below $25,000.[15]

[14]The lack of AOTC bunching is even more surprising given responsiveness to another education-related tax provision that was in effect prior to the introduction of the AOTC: a notch where tuition deductions are abruptly nullified (Hoxby and Bulman, 2016).

[15]In addition, in Appendix C we test for bunching at income eligibility thresholds for various transfer programs, including Medicaid, the Supplemental Nutrition Assistance Program, and disability benefits. We find none, though annual tax return data are not well suited to analyze these programs, whose eligibility criteria primarily depend on monthly income.

Our broad finding of unresponsiveness implies that taxpayers are insensitive to marginal tax rates in the neighborhood of most kinks. This could be driven by several mutually compatible causes. First, gathering information about the tax schedule is costly, and taxpayers may have imperfect knowledge of their local tax schedule (Chetty and Saez, 2013). Second, taxpayers may not base their decisions on marginal incentives (Feldman, Katuscak and Kawano, 2015). Third, taxpayers may know their local schedule and want to respond to marginal incentives, but may be constrained by optimization frictions such as adjustment costs or lumpy earnings opportunities. This explanation is consistent with Gelber, Jones and Sacks (2013), who analyze bunching in wage income by retirees, but it is unconvincing when deduction opportunities are present (e.g., at statutory kinks). Deductions, such as charitable giving, allow taxpayers to precisely manipulate their taxable income at the end of the year, after gross income is observed. Fourth, if income is sufficiently volatile, taxpayers may be unable to respond to kink points, especially if tax units have multiple income earners or income types.

As mentioned earlier, when kinks are small, we generally do not expect a bunching response. Chetty (2012) shows that at many kinks, ignoring the change in incentives leads to utility losses of less than 1% compared to utility-maximizing choice. In these cases, it may be rational for taxpayers to ignore marginal incentives, as the benefits likely do not exceed adjustment costs or the costs of acquiring information regarding tax incentives. In light of this, the lack of responsiveness at most middle-income and high-income kinks is unsurprising.

Taxpayers are not universally unresponsive, however. We observe bunching at several low- and middle-income kinks. Similar to patterns documented in Saez (2010) and Chetty, Friedman and Saez (2013), we find sharp bunching at the first EITC kink in all years of our sample. This is where the strongest bunching occurs. We also document bunching at the zeroth statutory kink, consistent with Saez (2010). In addition, we provide new evidence of bunching at the second EITC kink, CTC refundability plateau kink, MWPTC kink, and the second and third statutory kinks. Responsiveness at the two EITC kinks, the CTC kink, and the second statutory kink are displayed in Figure 2 for selected years and household types. We explore these patterns in detail in the following sections.

II.A Estimation Technique

When quantifying bunching behavior, the key issue is how taxpayers would behave in the absence of a kink. In particular, we must specify an alternative (local) tax schedule as well as the (local) distribution of income under the alternative tax schedule. We estimate this counterfactual behavior separately for two scenarios, corresponding to the two marginal tax rates (MTRs) that hold above and below the kink. We refer to the MTR that applies below the kink as t_0, and the MTR that applies above it as t_1. First, for those bunchers located below (left of) the kink, we estimate their behavior under a locally constant MTR equal to t_0. In other words, we assume their MTR continues unchanged throughout the kink region. Second, for those bunchers located above (right

Figure 2: Bunching at four kinks

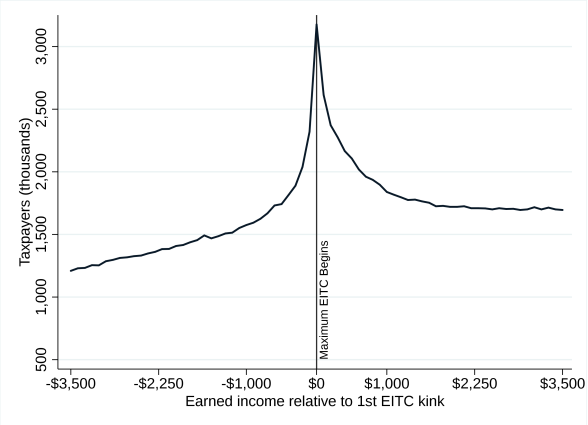

(a) 1st EITC kink (1996-2014)

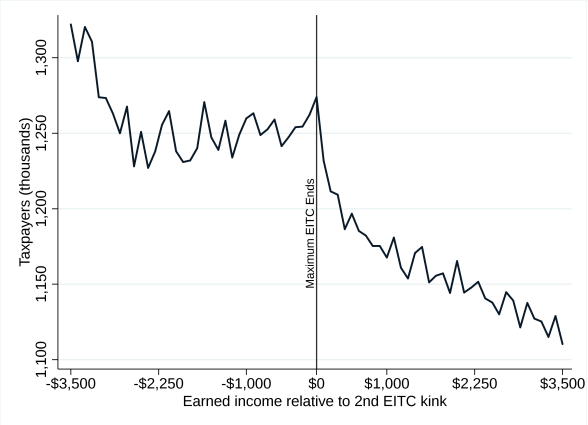

(b) 2nd EITC kink (2002-2014)

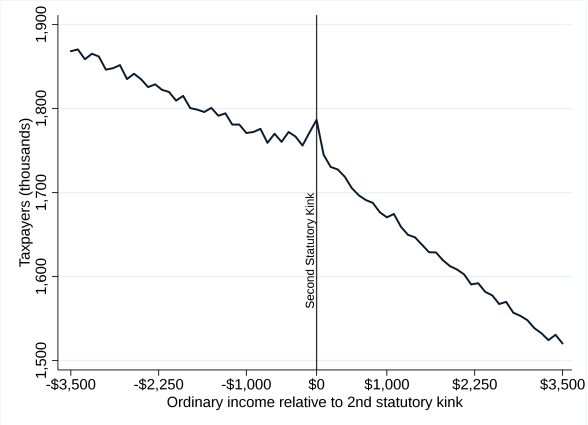

(c) 2nd statutory kink (1996-2014)

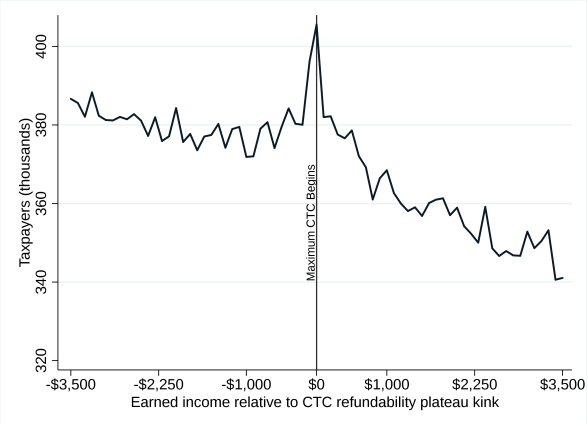

(d) CTC refundability plateau (2004-2014)

Panels (a) and (b) feature all EITC-eligible filers in our Main Sample, from 1996 to 2014 and 2002 to 2014, respectively, with the following exception. When the kinks are within $2,000, we drop all taxpayer types in (a) that respond to the second kink, and we drop all taxpayer types in (b) that respond to the first kink. Panel (c) includes all taxpayers, except those that have investment income, in all years of our sample. Panel (d) includes all taxpayers in our sample that have children, except those located within $2,000 of the first or second EITC kinks, from 2004 to 2014. Panels (a), (b), and (d) use $100 binwidths, while panel (c) uses a $200 binwidth. All panels were created by the authors using data from the IRS Compliance Data Warehouse and have been weighted to represent the full U.S. population.

of) the kink, we estimate their behavior under a locally constant MTR equal to t_1, assuming their MTR also held below the kink. Thus, the two counterfactuals are mutually incompatible.

In estimating these counterfactual scenarios separately, we break from the bunching analysis developed by Chetty et al. (2011). To our knowledge, all extant research that reports bunching coefficients uses their style of estimating one counterfactual distribution (under a constant MTR equal to t_0) for bunchers on both sides of the kink. We break from this approach for two reasons. First, their estimation equation is appropriate only for small kinks, as it is derived by taking the limit as t_1 converges to t_0. Second, Chetty et al. fit their counterfactual on the right of the kink to

the observed distribution of income under MTR t_1.[16] Saez (2010) shows that the distributions of income under MTRs t_0 and t_1 are generally not equal, nor are they directly proportional. Thus it is unclear what, if any, information the observed distribution above the kink offers when estimating the counterfactual distribution under MTR t_0.

For each counterfactual scenario, we estimate the income distribution using observed data near the kink but not so close as to be affected by bunching behavior. Specifically, we group households into bins and estimate distinct linear projections on both sides of the kink. For the counterfactual scenario where the MTR is t_0, we use bins $-R, ..., -1, 0$, where bin 0 contains the kink. For the counterfactual scenario where the MTR is t_1, we use bins $0, 1, ..., R$. We call the union of these sets of bins the "bunching region."

For the counterfactual scenario where the MTR is t_0, we estimate the following equation by ordinary least squares:

$$y_j = \alpha^0 + \beta^0 z_j + \sum_{k=-W}^{0} \gamma_k^0 \cdot \mathbf{1}[j = k] + \varepsilon_j^0, \tag{1}$$

where y_j denotes the number of taxpayers in bin j, z_j denotes the income level of bin j, W denotes the number of bins in the bunching window near the kink, and ε_j^0 denotes the residual.[17] Parameters γ_k^0 capture the number of taxpayers in the bunching window unexplained by the linear prediction $(\alpha^0 + \beta^0 z_k)$. In other words, γ_k^0 measures the amount of excess mass in bin k relative to the counterfactual expectation.

For the counterfactual scenario where the MTR is t_1, we estimate a similar equation:

$$y_j = \alpha^1 + \beta^1 z_j + \sum_{k=0}^{W} \gamma_k^1 \cdot \mathbf{1}[j = k] + \varepsilon_j^1. \tag{2}$$

Our default parameter values, which we select by visual inspection, are a bin-width of \$100 (δ=\$100), a bunching region of 71 bins (R=35), and a bunching window of 21 bins (W=10).[18] Our default counterfactuals are therefore derived from the actual distribution of income between \$1,000 and \$3,500 away from each kink. Letting circumflexes denote estimated coefficients, we calculate the total number of bunchers as $\hat{B} = \sum_{k=-W}^{-1} \hat{\gamma}_i^0 + \sum_{k=1}^{W} \hat{\gamma}_i^1 + (1/2)(\hat{\gamma}_0^0 + \hat{\gamma}_0^1)$. Figure 3 graphically depicts this estimation technique for married filers near the first EITC kink in 2005. The estimated number of bunchers is simply the difference between the observed and counterfactual distributions of income inside the bunching window.

In a few instances, two kinks are too close together to perform the analysis as described. Sup-

[16] Chetty et al. make an adjustment to their counterfactual, distributing the mass of bunchers uniformly to the right of the kink. Nonetheless, the general shape of their counterfactual on the right of the kink matches the shape of the observed distribution of income under MTR t_1.

[17] In practice, we normalize income levels z_j by de-meaning them and dividing by standard deviations. In exploratory work, we tried including higher-order polynomial terms of z_j, but this would generally over-fit the data, producing unrealistic counterfactual projections in the bunching window.

[18] We show that our results are robust to parameter choice in Appendix B.

Figure 3: Actual and estimated counterfactual distributions of income

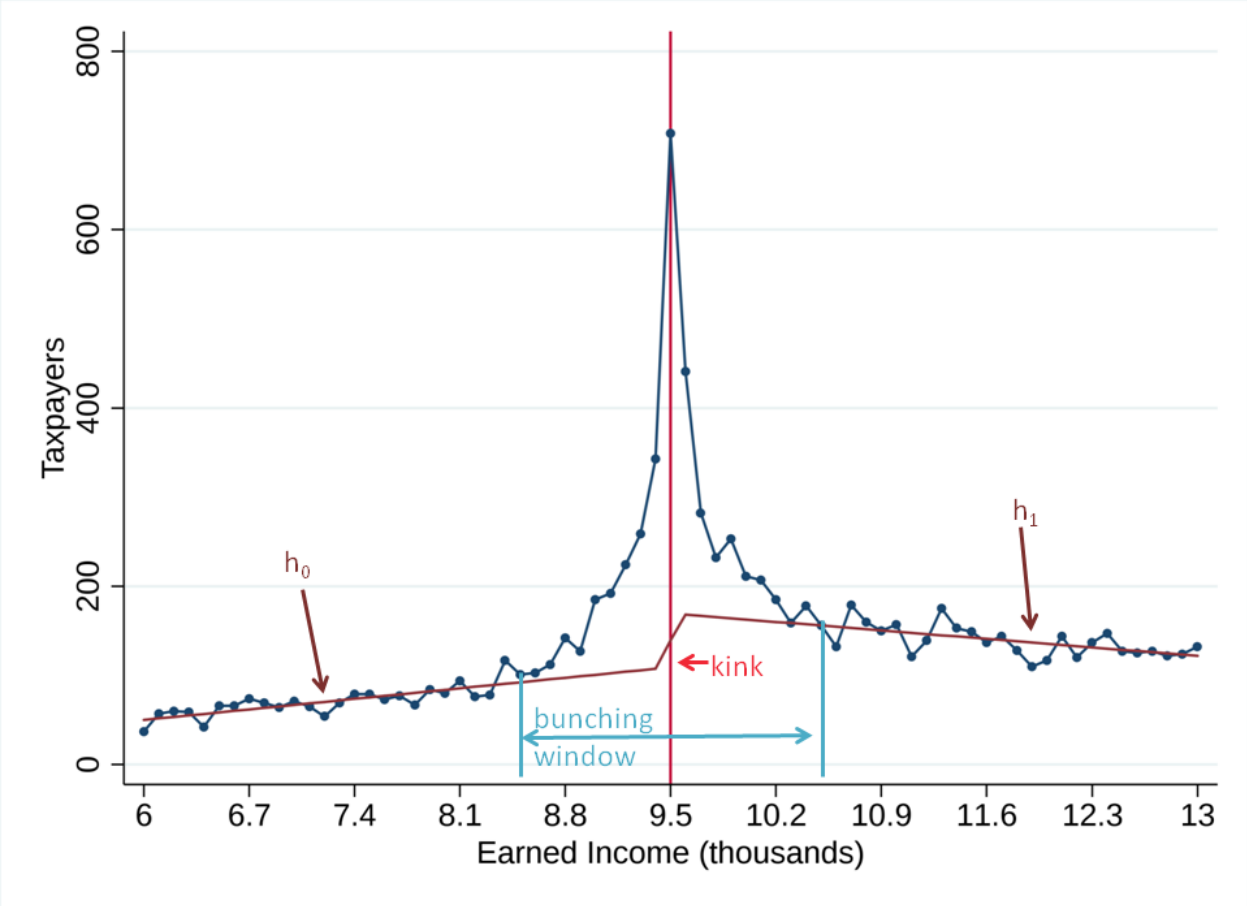

The distribution of income is displayed for married couples filing jointly in 2005 who have one dependent and who have self-employment income. The estimation parameters are the default ones: $R = 35$, $W = 10$, and $\delta = \$100$. Fitted lines h_0 and h_1 denote our estimates for the counterfactual densities of income under the assumptions that the constant marginal tax rate is equal to t_0 (the rate below the kink) and t_1 (the rate above the kink), respectively. This figure was created by the authors using data from the IRS Compliance Data Warehouse and is unweighted.

pose we wish to analyze kink K, but kink L lies somewhere inside K's bunching region. If taxpayers bunch at L, this can lead to unreasonable estimates for the counterfactual distributions needed for K's analysis. For this reason we do not report bunching coefficients for kink K whenever (i) taxpayers bunch at some kink L, and (ii) the distance between kinks K and L is between \$1,000 and \$2,000. When the distance between the kinks is less than \$1,000, so that kink L lies within kink K's bunching window, the problem is not the estimation of K's counterfactual distributions. Instead, the difficulty is that it is hard to tell which kink bunchers are responding to. In this case, we estimate the total number of bunchers in the usual way, except we divide them into two groups. If kink K sees marginal tax rates change by Δt_K, and kink L sees marginal tax rates change by Δt_L, then we assign fraction $\Delta t_K/(\Delta t_K + \Delta t_L)$ of the bunchers to K, and one minus this fraction to L.

Regardless of whether other kinks are nearby, the total number of bunchers is a flawed metric

for taxpayer responsiveness. All else equal, the number of bunchers will be larger when analyzing kinks affecting a larger mass of taxpayers. Hence, we report a bunching coefficient $\hat{b}$ equal to the percentage of taxpayers inside the bunching window who are classified as bunchers. In other words, letting P_k denote the observed population in bin k, we define $\hat{b} \equiv 100 \cdot \hat{B} / \sum_{k=-W}^{W} P_k$.[19] We use a bootstrap procedure to obtain standard errors for $\hat{B}$ and $\hat{b}$ by adding randomly sampled estimated residuals from the original regressions to the predicted values of the original regressions, repeatedly estimating $\hat{B}$ and $\hat{b}$ from the new, simulated data.[20]

II.B Bunching Estimation Results

Though all taxpayers face incentives to bunch at the convex kinks of Figure 2, some taxpayers are more responsive to these incentives than others. To compare bunching patterns across groups, Table 4 presents estimated bunching coefficients at four kinks where we find a response, using our most recent five years of data. In general, the first EITC kink elicits the largest response. It sees the largest bunching coefficient, 44.9%, corresponding to single, self-employed individuals. For this group, nearly half of all taxpayers within $1,000 of the first EITC kink during 2010 to 2014 were there because of the changing marginal incentives at the kink.[21]

Contrasting with prior research, we observe wage earners (i.e., those without self-employment income) bunching at many kinks in recent years. Table 4 shows that the average bunching coefficients for single wage earners at the first EITC kink and the CTC refundability plateau kink are more than four standard deviations away from zero. Though the table does not show it, single wage earners also exhibit statistically significant bunching coefficients at the second EITC kink and the second statutory kink in certain years. In contrast, married-filing-jointly wage earners do not exhibit statistically significant coefficients in Table 4 or, in fact, in any year.[22]

In general, bunching by wage earners is smaller than bunching by those with self-employment income. All of the statistically-significant bunching coefficients for the self-employed are larger than those of their wage-earning counterparts in Table 4, and most of these differences are themselves statistically significant. However, it is unclear whether the greater responsiveness of the self-employed can be attributed to labor supply decisions or higher rates of tax evasion. On the one hand, adjusting earned income is inherently easier when one is both the employer and employee,

[19] When analyzing kink K, if kink L is nearby, we adjust the denominator in the definition of K's bunching coefficient, $\hat{b}$, by subtracting the bunchers $\hat{B}$ attributed to kink L.

[20] We thank Raj Chetty, John Friedman, Tore Olsen, and Luigi Pistaferri for public provision of a Stata program designed specifically to implement their estimation technique. Our code builds directly on theirs, and we plan to make our code publicly available in the near future.

[21] Other groups and other kinks exhibit smaller, statistically insignificant bunching coefficients that are occasionally negative. Negative numbers imply the kink causes *less* mass to locate near the kink. This is plausible only taxpayers respond to higher tax rates by earning *more* income, which has little empirical support. As none of the negative coefficients are statistically distinguishable from zero, we interpret them as evidence that taxpayers are not responding to the kink.

[22] Married-filing-*separately* wage earners, however, bunch at the third statutory kink in about half of the sample years (see Appendix A).

Table 4: Average bunching coefficients at four kinks (2010-2014)

	1st EITC kink	2nd EITC kink	CTC refundability plateau	2nd statutory kink
Single wage earners	5.0% (1.2%) [N=1,345,000]	-0.5% (1.3%) [N=1,536,000]	4.2% (0.6%) [N=1,252,000]	0.7% (0.5%) [N=1,919,000]
Single self-employed	44.9% (1.4%) [N=909,000]	3.8% (3.9%) [N=657,000]	29.8% (1.0%) [N=482,000]	1.3% (1.9%) [N=107,000]
Married-filing-jointly wage earners	2.3% (3.9%) [N=298,000]	-0.4% (2.9%) [N=493,000]	1.6% (2.3%) [N=332,000]	0.5% (0.5%) [N=976,000]
Married-filing-jointly self-employed	19.3% (2.7%) [N=235,000]	3.6% (2.8%) [N=210,000]	16.2% (1.8%) [N=207,000]	1.0% (1.4%) [N=169,000]
Married-filing-sep. wage earners	—	—	1.1% (9.8%) [N=15,000]	1.6% (1.8%) [N=110,000]
Married-filing-sep. self-employed	—	—	13.1% (11.3%) [N=2,000]	7.2% (6.9%) [N=8,000]

Weighted-average bunching coefficients – the percentage of taxpayers within the bunching window that are estimated to be bunching – are reported for various household types, with standard errors in parentheses. In all four columns, the bunching window includes taxpayers within $1,000 on either side of the kink. The number of taxpayers in the bunching region (rounded to the nearest thousand) is presented in brackets. Wage earners are those with positive wage income and zero self-employment income. The self-employed are those with nonzero self-employment income. Single status includes "head of household" filers. Estimates are omitted when the kink is between $1,000 and $2,000 away from another kink where taxpayers bunch, as discussed in Section II.A. Married-filing-separately taxpayers are ineligible for the EITC and are excluded from its analysis. The income definition is earned income for the EITC and CTC kinks and taxable income for the statutory kink. All figures reflect the authors' calculations using data from the IRS Compliance Data Warehouse.

so the self-employed likely exhibit larger real labor responses. Moreover, self-employment income is easily adjusted at the end of the tax year by incurring business expenses. On the other hand, unlike wage income, self-employment income is not subject to third-party reporting and is therefore easier to hide from tax authorities (or inflate). Thus, the self-employed likely exhibit greater tax evasion responses as well.

Unlike bunching by the self-employed, bunching by wage earners provides prima facie evidence of a real labor response to taxation. Because wage earnings are reported by third parties, they are thought to be more reliable indicators of real economic activity (Slemrod, 2007). Indeed, the Internal Revenue Service estimates that 99% of aggregate taxable wage income is reported by taxpayers on Form 1040 (Internal Revenue Service, 2016). Further, business expenses and itemized deductions are irrelevant for wage earners bunching at EITC and CTC kinks, which are determined by gross income alone. If bunching behavior in response to these credits represents a distortion in the labor market, rather than simply a reporting phenomenon, this would change

the welfare consequences of the credits (Chetty, 2009). However, in Section III.A we show that bunching in taxpayer-reported wages does not manifest in employer-reported wages. Instead, the wage-earner bunching we uncover merely reflects misreporting of wages by taxpayers on Form 1040 relative to wages reported on Form W-2.

In addition to the four kinks of Table 4, in Appendix A we document weak, but statistically significant, bunching at the zeroth and third kinks in the statutory schedule. The former coincides with the filing threshold for wage earners, and therefore suffers from censoring in the data. In addition, taxpayers with children do not see a change in marginal incentives at this kink because the non-refundable portion of the Child Tax Credit immediately eliminates their liability. Thus, we can cleanly measure bunching here only for childless, self-employed taxpayers. The bunching coefficients we observe are more stable over time for married taxpayers than singles, but both groups produce small, statistically significant bunching in roughly half of the years in the sample.

At the third statutory kink we face no inherent sample restrictions, yet we find statistically significant bunching only for married-filing-separately taxpayers. Other groups, including married couples filing jointly, fail to respond during any years of our sample. At this kink and all kinks above it, we analyze the universe of tax returns because our 10% sample yields insufficient mass to distinguish bunching from noise. 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.

Finally, we also find responsiveness to the temporary Making Work Pay Tax Credit by married, self-employed taxpayers with two or more children. In 2009 and 2010, this credit created a convex kink of roughly six percentage points at the end of its phase-in. When fully phased in, the credit delivered $800 to married taxpayers filing jointly. Unfortunately, given its proximity to the first EITC kink for responsive households, our bunching estimates require assumptions about which bunchers near the kink are assigned to the EITC and MWPTC kinks. Nonetheless, Figure 4 clearly shows separate responses to each kink, indicating some taxpayers were responding specifically to the MWPTC. This is somewhat surprising, as the MWPTC kink changed marginal tax rates by six percentage points, whereas the first EITC kink changed rates by 40 percentage points for those with two children and 45 percentage points for those with three children. Standard models of taxpayer behavior, which predict bunching responses proportional to kink size, are difficult to reconcile with the observed distribution peaking at the smaller MWPTC kink. However, for the vast majority of taxpayers included in Figure 4, the MWPTC kink marked the unique spot in the income distribution that maximized their refunds. In Section III.B, we show that gravitation towards the refund-maximizing kink is not unique to this group.

II.B.1 Evolution of bunching patterns

One of the most striking features of the bunching patterns we observe is their evolution over time. In 1996, substantial bunching occurred only at the first EITC kink and only for those with self-

Figure 4: Responses to the Making Work Pay Tax Credit (2010)

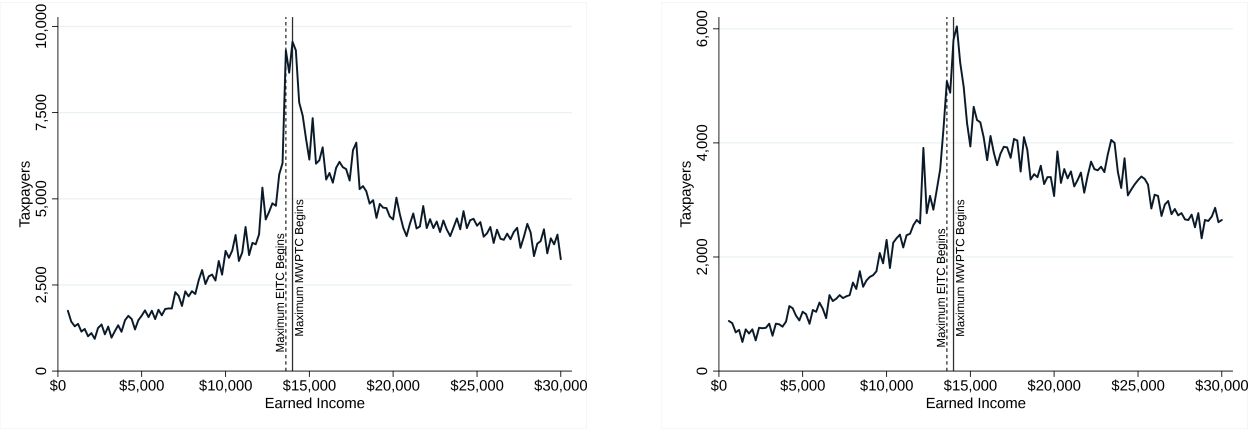

(a) Married, self-employed, 2 children

(b) Married, self-employed, 3+ children

Income distributions are displayed for married, self-employed taxpayers with two or more children. The dashed vertical line denotes the first EITC kink; the solid line denotes the MWPTC kink. The self-employed are those with nonzero self-employment income. Married taxpayers who file separately from their spouse are excluded. The figures look very similar using 2009 data. All panels were created by the authors using data from the IRS Compliance Data Warehouse and have been weighted to represent the full U.S. population.

employment income. We estimate that the total number of bunchers at the kinks we study was 104,000. By 2014, there is substantial bunching at both the first EITC kink and the CTC refundability plateau, including wage-earner bunching. The total number of bunchers had risen by 700% to 834,000, dramatically outpacing the 20% rise in tax returns filed over the same period.

Figure 5 displays this temporal variation. Several findings emerge from studying the figure. First, both panels indicate bunching patterns are not monotonically increasing over time. Second, panel (b) shows that substantial wage-earner bunching is a relatively recent phenomenon. This is why we find wage-earner bunching where Saez (2010) and Chetty, Friedman and Saez (2013) found none, as their samples end in 2004 and 2009, respectively. Third, panel (a) shows that bunching in response to the CTC emerges in 2009, with the total number of CTC bunchers rapidly ascending to levels previously seen only at the first EITC kink. This is especially noteworthy given that the CTC kink changes effective marginal tax rates by 15 percentage points, whereas the first EITC kink sees effective marginal tax rates rise by up to 45 percentage points. Fourth, bunching at the first EITC kink (and, consequently, overall bunching) drops in 2008 at the onset of the Great Recession, as depicted in panel (a). An intriguing hypothesis is that economic downturns may influence taxpayers' ability to bunch, consistent with evidence from Ireland during the Great Recession (Hargarden, 2015). If this conjecture holds, it would provide evidence against the "income reporting" explanation for bunching, suggesting real economic activity may be involved.

Panel (b) depicts a general upward trend in bunching at refund-maximizing kinks. We estimate that the number of bunchers locating at the refund-maximizing kink rose from 90,000 to 616,000 between 1996 and 2014. However, in a few instances refund-maximizing bunching falls quite dramatically. For example, the years 2011 and 2012 stand out as particularly off-trend. As we

Figure 5: Bunching over time

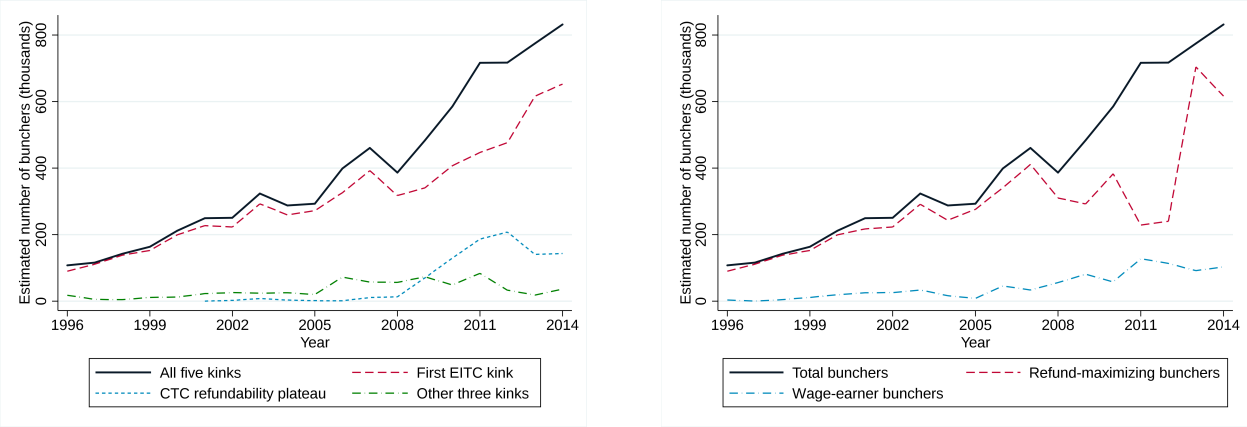

(a) Bunchers at specific kinks

(b) Wage earners and refund maximizers

Estimates for the total number of bunchers at five kinks (the first and second EITC kinks, the CTC refundability plateau kink, the MWPTC kink, and the second statutory kink) are displayed. For a given group at a given kink in a given year, we impose that the number of bunchers is zero whenever the estimated bunching coefficient is less than two times its standard deviation (including instances when the coefficient is negative). Refund-maximizing kinks are identified using NBER TAXSIM. These figures were created by the authors using data from the IRS Compliance Data Warehouse and have been weighted to represent the full U.S. population.

will see in Figures 7 and 8, many taxpayers saw the refund-maximizing kink switch from the first EITC kink to the CTC refundability plateau in 2011. While some taxpayers followed the refund-maximizing kink as it changed location during this period, many others did not. By 2013, however, the first EITC kink was once again the refund-maximizing kink for the vast majority of taxpayers. Accordingly, in 2013 two things occur: the number of refund-maximizers returns to its pre-2011 trend, and, in panel (a), bunching at the CTC kink falls.

III Characteristics of Bunchers

Thus far we have focused on the intensity and location of bunching patterns. In this section, we explore *who* bunches, *how* they bunch, and *why* different groups are drawn to different kinks. We present three key facts about taxpayer bunching. First, we show that wage-earner bunching is merely a reporting phenomenon, rather than evidence of labor market distortions. Second, we demonstrate that bunching responses are significantly more likely to occur at the unique point in the schedule that results in the largest refund, maximizing tax credits net of taxes owed. Third, we show that taxpayer bunching is more than a transitory phenomenon: taxpayers persist at kinks year after year, especially at the first EITC kink. We begin with graphical evidence of these behaviors among specific groups. We conclude with regression analysis, showing the patterns hold in the full sample as well.

III.A Wage-Earner Bunching: Real Response or Reporting?

The bunching responses described in Section II imply that taxpayers (or paid preparers) have a sophisticated understanding of the tax code as well as the ability to precisely control their reported incomes. How are taxpayers able to achieve this? In this section we investigate the mechanisms behind taxpayer bunching.

Unfortunately, when taxpayers are self-employed it is difficult to disentangle the various mechanisms available for bunching. As these taxpayers are both employer and employee, labor supply and demand responses are conflated. And because the tax schedule is defined with respect to *net* self-employment income, it is hard to distinguish income responses from deduction responses. If a taxpayer reports $15,000 in income and $5,000 in expenses in order to bunch at a $10,000 kink, we cannot know whether she first chose income and then adjusted expenses to land near the kink, or vice versa. Further, because net self-employment income is not subject to third-party reporting, it is difficult to detect non-compliance without thorough audits.[23]

We have more traction, however, with analysis of pure wage earners (i.e., those without self-employment income). For this group, the relevant income concept at the kinks where we observe large responses to the EITC and CTC is gross earned income. This does not allow for deductions, immediately ruling out this response channel. In addition, all employers that pay an employee more than more than $600 in a given year are required to file Form W-2 with the IRS.[24] Due to this third-party reporting, wage income has been shown to be subject to significantly less non-compliance than self-employment income (Slemrod, 2007; Internal Revenue Service, 2016). Thus, wage-earner bunching provides prima facie evidence of real labor supply or real labor demand responses. However, because we observe Form W-2, we can detect noncompliance directly by comparing taxpayer- and employer-reported incomes. The former comes directly from line 7 of Form 1040. For the latter, we sum wage income across all Forms W-2 filed for a given taxpayer (or his spouse) and add to this any additional wage or tip income reported on Forms 4137 or 8919.[25]

Our results indicate that the excess mass of wage earners near kinks is wholly due to misreported earnings. Panels (a) and (b) of Figure 6 display the distributions of taxpayer- and employer-reported wage income for taxpayers with two children in 2014. The patterns are broadly similar for those with different numbers of children in other recent years. Relative to employer-reported income, taxpayer-reported income in features additional mass in the EITC plateau region (between the two solid vertical bars) for the single taxpayers of panel (a). This extra mass exhibits sharp bunching precisely at the CTC refundability plateau (marked by the dashed bar). Employer-reported wages in panel (a), however, show no indication of bunching at any kink. This rules out

[23] In a tax audit experiment in Denmark, Kleven et al. (2011) find that around half of the bunching response of the self-employed is eliminated post-audit. We intend to explore this in the U.S. context in future research.

[24] If any taxes are withheld, even employees earning less than $600 must have a W-2 filed. Household employees (e.g., nannies) have a looser threshold: $1,900 in 2014.

[25] Any wage or tip income not reported by the employer on Form W-2 must be reported by the taxpayer on Forms 4137 and 8919. Thus, any mismatch between our employer- and taxpayer-reported income measures necessarily reflects noncompliance.

Figure 6: Taxpayer- and employer-reported wage income (2014)

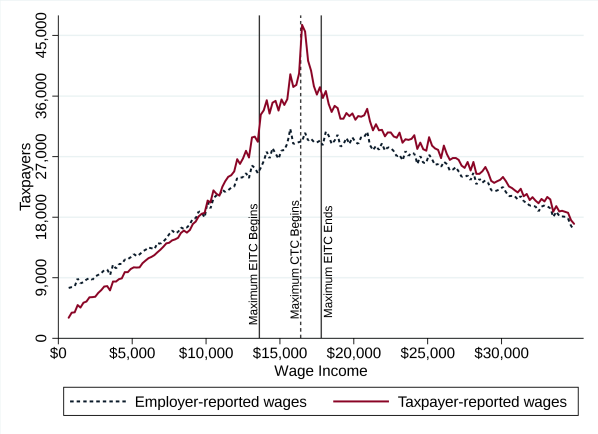

(a) Single, wage earners, two children

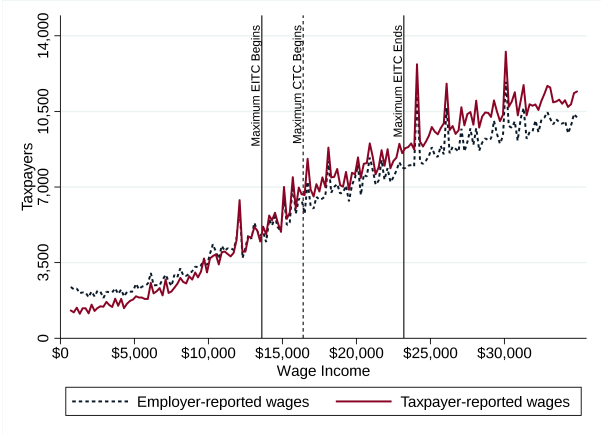

(b) Married, wage earners, two children

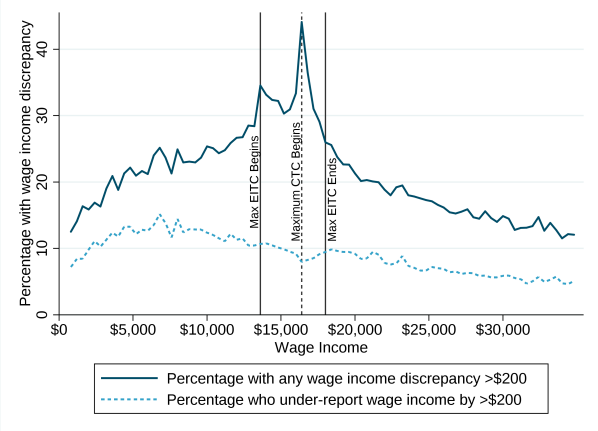

(c) Single, wage earners, two children

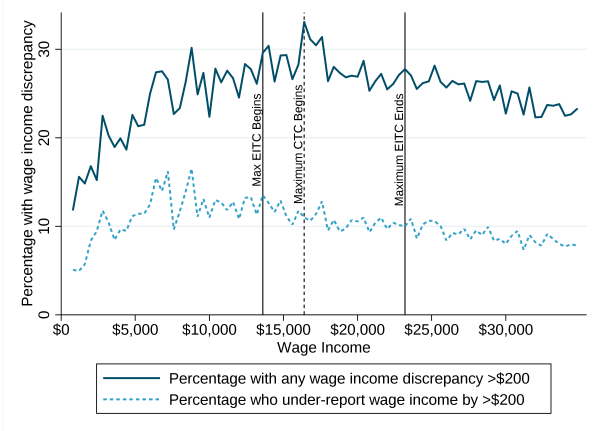

(d) Married, wage earners, two children

Distributions of taxpayer-reported and employer-reported wages from $600 to $35,000 are displayed in panels (a) and (b) for single and married taxpayers, respectively. Panels (c) and (d) show the percentage of taxpayers who report income in a certain bin who have a mismatch between their taxpayer- and employer-reported incomes of more than $200. The solid vertical bars mark the first and second EITC kinks, and the dashed vertical bar denotes the CTC refundability plateau. Wage earners are those with positive wage income and zero self-employment income. Single status includes "head of household" filers. Married taxpayers who file separately are excluded. Taxpayer-reported wages taken from line 7 of Form 1040, whereas employer-reported wages are taken from Form W-2 and include additional wage and tip income reported on Forms 4137 and 8919. All panels were created by the authors using data from the IRS Compliance Data Warehouse. Panels (a) and (b) have been weighted to represent the full U.S. population.

the possibility of a real labor supply or labor demand response to kinks, leaving non-compliance as the culprit.[26]

The evidence is different for the married taxpayers of panel (b). There is little excess mass in the EITC plateau region, and no bunching in either taxpayer-reported or employer-reported wages. In both panels (a) and (b), we see more mass in taxpayer-reported wages above $15,000. However, in both cases the two lines spike in parallel, suggesting this mismatch is random – perhaps due to

[26]Some bunching may be due to identify theft, wherein a criminal obtains a Social Security number and files a fraudulent return, attempting to collect the victim's refund. To the extent this behavior is detected by the IRS at the time of filing, however, such returns are not included in our data.

imperfect matching of tax returns with Forms W-2.

Panels (c) and (d) of Figure 6 corroborate the evidence in panels (a) and (b). The panels display the percentage of taxpayers in a given income bin who have a discrepancy between their employer- and taxpayer-reported incomes greater than $200. For married individuals, there is little evidence of systematic mismatch, save perhaps at the CTC refundability plateau. For singles, however, the evidence is clear that taxpayers reporting incomes in the bins containing the first EITC kink and the CTC refundability plateau are more likely to have a wage-income discrepancy. Nearly 45% of taxpayers in the bin containing the CTC refundability plateau have a discrepancy, whereas outside the EITC plateau region the percentages do not exceed 30 and typically fall between 10 and 25.

Panels (c) and (d) also indicate that wage income discrepancies are not driven by under-reporting. The percentage of taxpayers who under-report wage income does not spike at the kinks. This means bunching due to wage income discrepancies is driven by taxpayers reporting incomes exceeding those their employers report. This is consistent with incentives because when incomes are subsidized – as they are below the CTC kink in this case – fabrication of income decreases tax liability, increasing refunds. In Figure 6, taxpayers are particularly likely to over-report wage income in order to bunch at the CTC kink, which marks the unique refund-maximizing kink for 98% of the taxpayers depicted. We explore this strategic reporting behavior in further detail in the next section.

Though we have shown that some taxpayers misreport wage income to bunch at kinks, the success rate of this strategy is unclear. The IRS expends considerable resources auditing taxpayers who claim refundable tax credits, as their returns are known to commonly contain errors.[27] Because our data do not contain post-audit outcomes, we do not know whether the refunds claimed by taxpayers are ever paid out or are recovered later through audit procedures. As wage-income discrepancies are easily identified, we suspect the IRS likely detects a large share of this noncompliance.

III.B Bunching to Maximize Refunds

In Section II.B.1, we demonstrated that bunching patterns are not constant over time, with changes at both the extensive and intensive margins. For example, bunching at the CTC refundability plateau and the second EITC kink is nonexistent until emerging in the 2000s. Similarly, for certain groups bunching coefficients rise at one kink at the same time another kink sees coefficients fall. We hypothesize that this is due to taxpayers seeking the point in the schedule that maximizes tax credits net of taxes owed, earning the maximum possible refund (given fixed withholding decisions).

Figure 7 shows this refund-maximizing behavior in action.[28] Here we see the evolution of

[27] In a study of random audits, the IRS estimates that wage income misreporting led taxpayers to over-claim the EITC by around $0.8 to $1.1 billion (in 2008 dollars) during 2006-2008 (Internal Revenue Service, 2014).

[28] We calculate the location of refund-maximizing kinks using NBER's TAXSIM software (Feenberg and Coutts, 1993), available at `http://www.nber.org/taxsim/`.

Figure 7: Tracking the kink that maximizes credits net of taxes owed: Singles with two children and with self-employment income

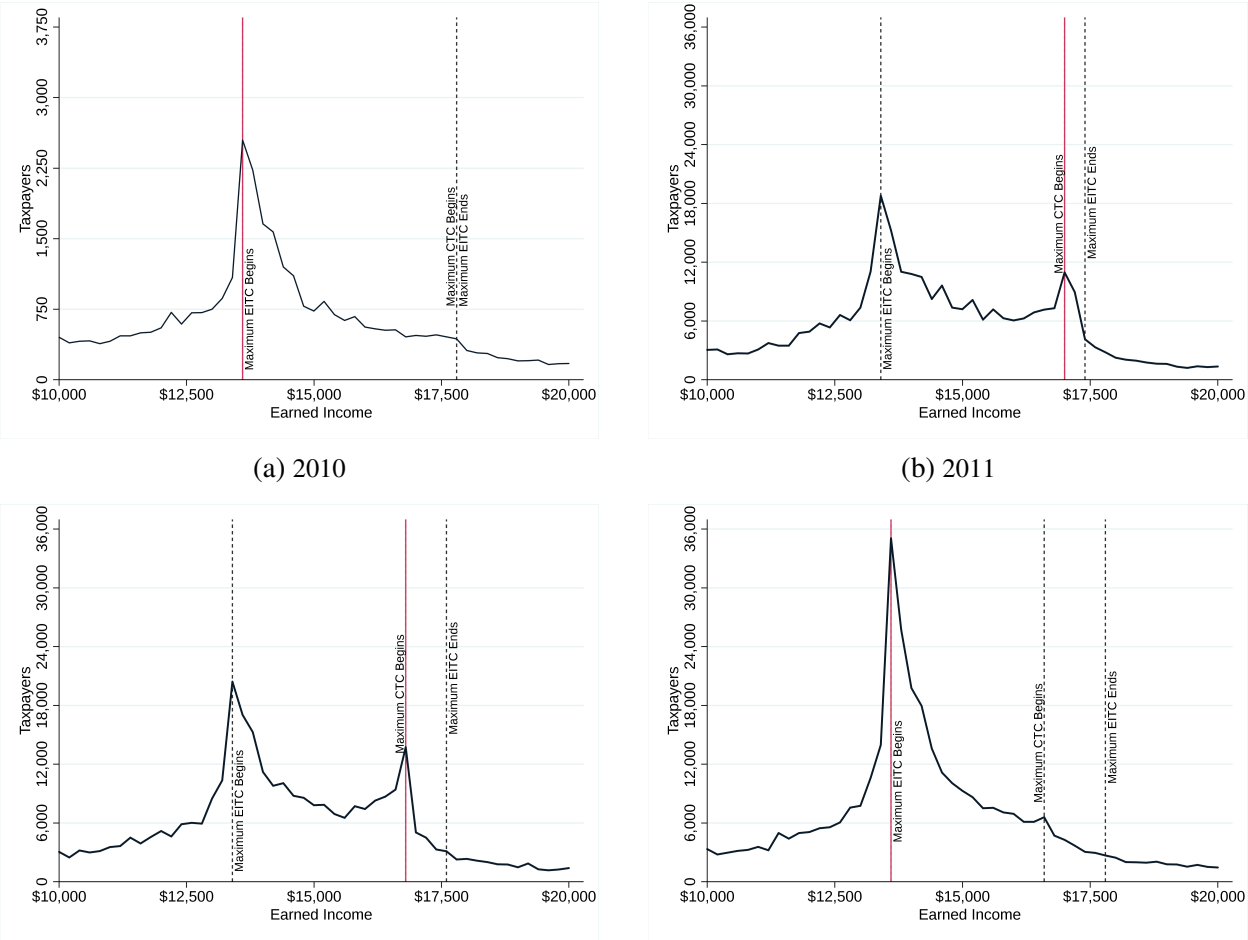

Annual income distributions are displayed for single taxpayers with two children and positive self-employment income in 2010-2013. The sample is limited to those taxpayers whose refund-maximizing kinks moved from the first EITC kink in 2010 to the CTC refundability kink in 2011, and from the CTC refundability kink in 2012 back to the first EITC kink in 2013. The solid vertical line denotes the kink that maximizes refunds, while dashed vertical lines denote other kinks where taxpayers respond. Single status includes "head-of-household" filers. All panels were created by the authors using data from the IRS Compliance Data Warehouse and have been weighted to represent the full U.S. population.

the income distribution for self-employed singles with two dependents from 2010 to 2013. In all panels, the point in the schedule that results in the maximum refund is marked by a solid (red) vertical line. In 2010 this point is the first EITC kink for this group. However, in 2011 the refund-maximizing point shifted to the CTC refundability plateau. This happened for two reasons. First, the CTC kink, which is not indexed to inflation, dropped below the second EITC kink, which is indexed to inflation. Second, a temporary two percentage point reduction in employee payroll taxes was enacted in 2011 as part of the Tax Relief, Unemployment Insurance Reauthorization, and Job Creation Act of 2010. This changed effective marginal tax rates above the first EITC kink from slightly positive to slightly negative.

Figure 8: Tracking the refund-maximizing kink: Married couples with two children and with self-employment income

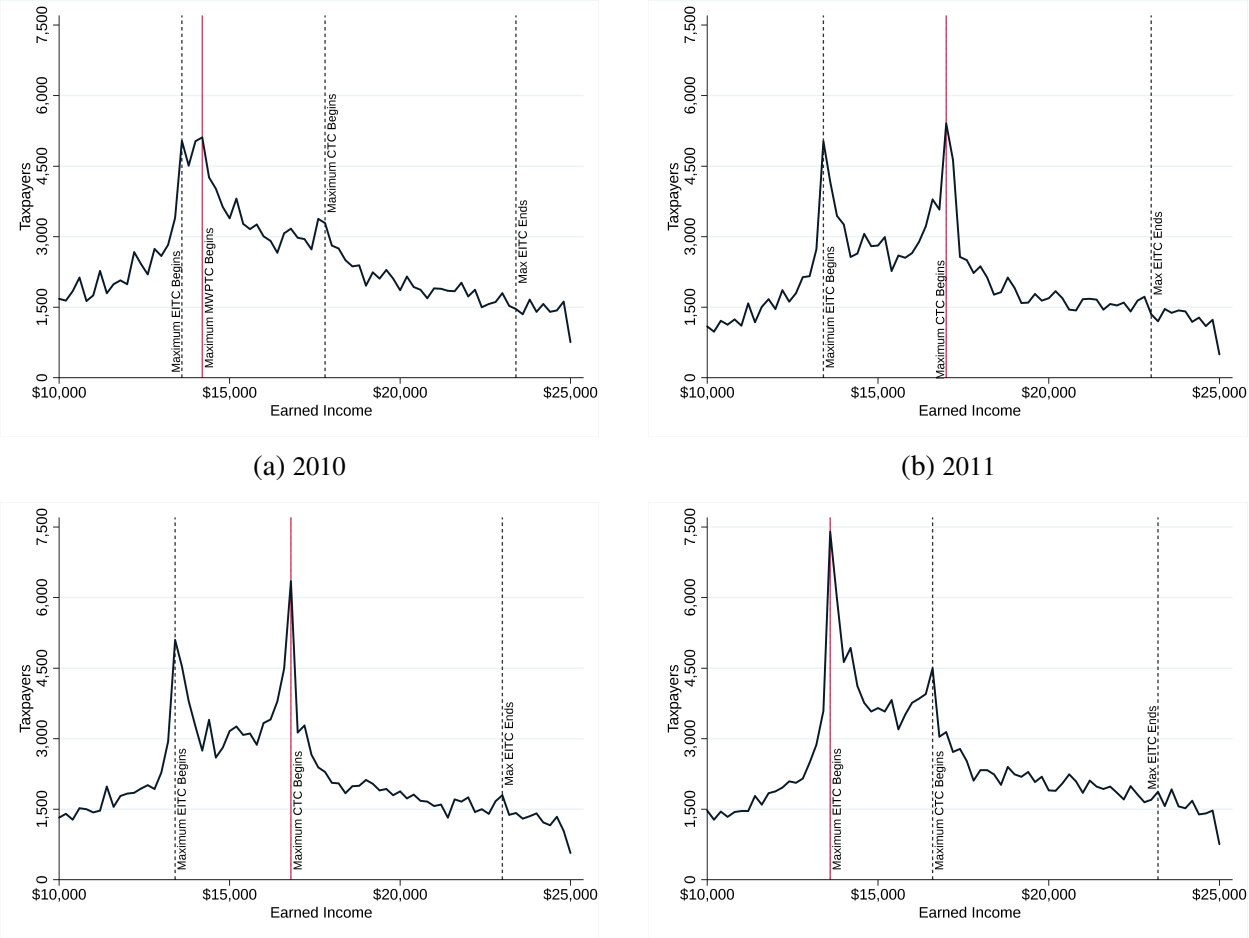

(a) 2010

(b) 2011

(c) 2012

(d) 2013

Annual income distributions are displayed for married taxpayers with two children and positive self-employment income in 2010-2013. The sample is limited to those taxpayers whose refund-maximizing kinks moved from the first MWPTC kink in 2010 to the CTC refundability kink in 2011, and from the CTC refundability kink in 2012 back to the first EITC kink in 2013. The solid vertical line denotes the kink that maximizes refunds, while dashed vertical lines denote other kinks where taxpayers respond. Married-filing-separately taxpayers are excluded. All panels were created by the authors using data from the IRS Compliance Data Warehouse and have been weighted to represent the full U.S. population.

Evincing their sophistication, taxpayers immediately respond to this minor change in the tax code. Relative to 2010, the 2011 income distribution features significantly less bunching at the first EITC kink and significantly more at the CTC kink. The following year, the distribution of income looks similar, as the CTC kink remains the point in the schedule that maximizes refunds. Then, in 2013, the payroll tax holiday expires and the first EITC kink reclaims its status as the refund-maximizing kink. Once again, taxpayers immediately respond by shifting to the new optimum; bunching plummets at the CTC kink while it soars at the first EITC kink.

We see this behavior among married couples as well. Figure 8 displays income distributions for the married counterparts to Figure 7. These taxpayers also see the refund-maximizing point in

the schedule shift from the first EITC kink to the CTC refundability plateau in 2011 and then back to the first EITC kink in 2013. Like their single counterparts, they respond immediately to each of these shifts.

Given that the size of each kink remains essentially constant during the periods highlighted in Figures 7 and 8, it is difficult to square the evidence with standard models of labor supply. Typically, we expect taxpayers to respond to marginal incentives given fixed preferences, which would suggest that the amount of bunching at a given kink would be proportional to its size. This assumption underlies the estimation technique that translates bunching patterns into elasticities of taxable income (Saez, 2010; Chetty et al., 2011). We do not report elasticities in this paper, precisely because of the evidence in Figures 7 and 8. If taxpayers merely report incomes to maximize refunds, standard methods to calculate elasticities do not apply. This is because kink size – which corresponds to the denominator in the elasticity – is irrelevant to the income reporting decision.[29] Indeed, the only aspect of the tax schedule that matters to a refund-maximizer is the point where marginal rates flip from negative to positive.[30]

III.C Learning to Bunch? Examining Bunching Persistence

We have seen that, for some groups of taxpayers, income distributions shift over time to track refund-maximizing kinks. This suggests a dynamic element to bunching: some taxpayers may repeatedly report incomes that maximize tax credits. In this section, we explore this behavior in further detail, measuring persistence at kinks when the refund-maximizing kink does and does not change locations. Because this analysis relies on tracking individuals over time, we use our Combined Sample, described in Section I.A, which lets us track taxpayers who alternate between being the primary and secondary filer in a tax unit. We limit analysis to those in our Main Sample in the base year, ensuring our results remain nationally representative.

We measure persistence in incomes (i.e., the absence of income volatility), comparing taxpayers near kinks to those away from kinks. Our measure for persistence is an indicator variable for remaining within $500 (in 2014 dollars) in real income in consecutive years. Figure 9 displays these persistence measures for self-employed individuals with one or two children, using 2011 as the base year. For all groups depicted, the refund-maximizing kink was the CTC refundability plateau in 2011 and 2012, but switched to the first EITC kink in 2013. For those with one child, the two kinks are within $500 during 2011-2013, and coincide in 2013. For these taxpayers, depicted

[29] In an ideal research design, one could calculate elasticities from bunching measures in the standard way by identifying refund-maximizers and removing them from the data. Alternatively, one could simply estimate elasticities at kinks other than refund-maximizing kinks. We refrain from this, too, as we suspect many bunchers at non–refund-maximizing kinks may be naively maximizing one tax credit (e.g., the EITC) as a heuristical approximation to refund maximization.

[30] A simple theoretical model that results in such targeted reporting behavior can be found in Section II.C of Saez (2010); however, the elasticities Saez reports are incompatible with this model. Alternatively, one can conceptualize refund-maximization as the solution to a variant of the model of Allingham and Sandmo (1972) in which: (i) taxpayers believe the probability of an audit is zero, and (ii) taxpayers may over-report incomes.

Figure 9: Bunching persistence (2011)

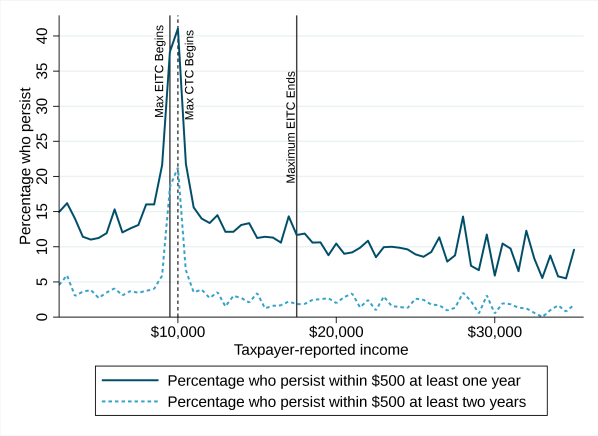

(a) Single, self-employed, one child

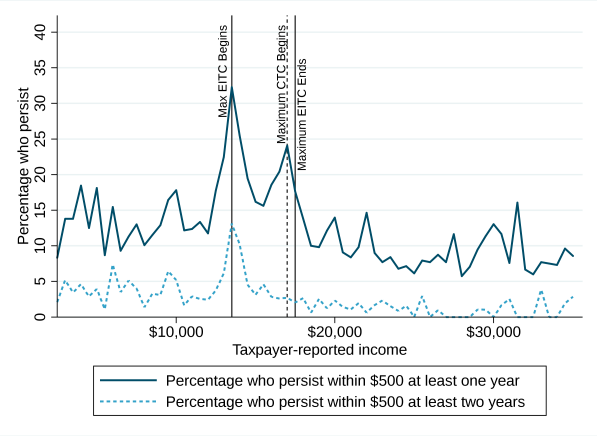

(b) Single, self-employed, two children

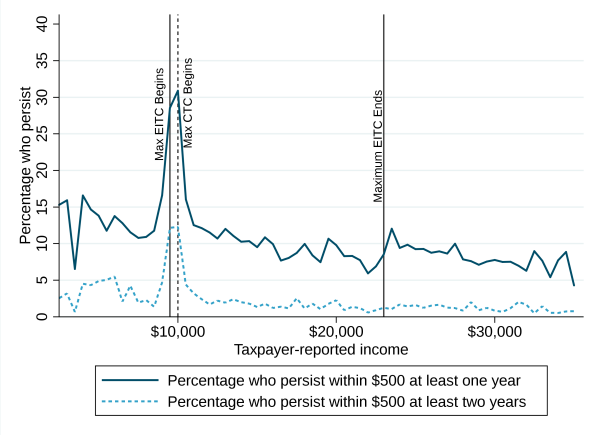

(c) Married, self-employed, one child

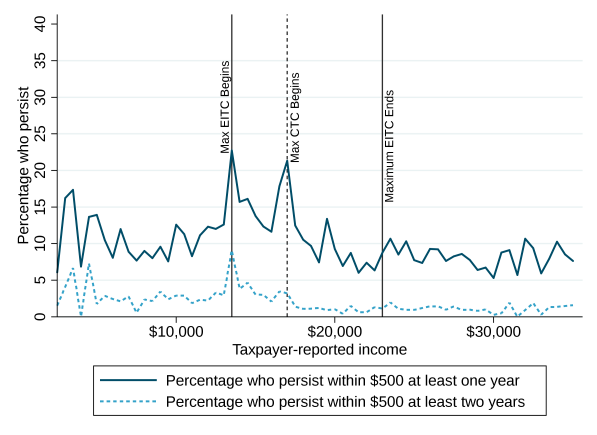

(d) Married, self-employed, two children

The graph displays the percentage of taxpayers in income bins of width $500 who persist within $500 of their current income, limiting the sample to self-employed taxpayers with one or two children in 2011. The solid vertical bars mark the first and second EITC kinks, and the dashed vertical bar denotes the CTC refundability plateau. Wage earners are those with positive wage income and zero self-employment income. Single status includes "head of household" filers. Married taxpayers who file separately are excluded. All dollar amounts are measured in 2014 dollars. All panels were created by the authors using data from the IRS Compliance Data Warehouse.

in panels (a) and (c), persistence near the kinks is significantly higher than persistence away from the kinks. Those near the first EITC kink or CTC kink in 2011 exhibit roughly a 25% chance of reporting similar incomes the following year and about a 9% chance of reporting similar incomes the following two years. Comparable figures away from the kinks, including those at the second EITC kink, are about 10% and 2%, respectively.

Taxpayers with two children, depicted in panels (b) and (d), provide even more insight. They, too, show elevated levels of persistence near the first EITC kink and the CTC kink. However, for those near the CTC kink, only the one-year measure of persistence shows a clear spike. For the singles of panel (b), the two-year measure of persistence indicates that those bunchers who persist one year at the CTC kink (in 2012) do not remain there a second time (in 2013). This is consistent

with refund-maximization, as this group sees the refund-maximizing kink move by more than $500 in 2013 (from the CTC kink to the first EITC kink). The same story holds for the married taxpayers of panel (d).

However, the pursuit of maximum refunds cannot fully explain the persistence results. In both panels (b) and (d), the largest spikes in persistence occur at the first EITC kink. This suggests many taxpayers simply target the maximum EITC, ignoring the incentives created by the CTC. One possible explanation for this is greater awareness about EITC incentives, as the EITC is the more established transfer program, featured fewer changes in incentives during the years leading up to 2011, and offers larger credit amounts than the CTC.

III.D Regression Evidence

The previous sections offered graphical evidence of income misreporting, refund maximization, and persistence in bunching among specific groups in selected years. Here we confirm that these patterns hold in the full sample as well. We begin by testing whether certain characteristics, such as a mismatch between taxpayer- and employer-reported incomes, are correlated with reporting incomes near a kink. Inevitably, the coefficients we obtain are attenuated towards zero because some non-bunchers locate near kinks by chance. Nonetheless, if a given variable is meaningfully correlated with the decision to bunch, we should obtain significant correlations in the regressions below.

Table 5 displays the results of probit regression analyses. Each column represents a different kink, and in every case the dependent variable is an indicator for whether the taxpayer is within $500 (in 2014 dollars) of the kink in question. A battery of dummy variable covariates are included, detailed in the note below the table, and errors are clustered at the taxpayer level. The coefficients are average marginal effects measured as percentage point increases in the likelihood of locating near that column's kink. In each regression, we limit the analysis to taxpayers with earned income within $3,500 (in 2014 dollars) of that column's kink. The EITC and refund-maximum regressions are further limited to EITC-eligible taxpayers.

The results indicate that over-reporting wages is positively correlated with locating near all four kinks. The largest effect is at the CTC refundability kink, where over-reporting wages is associated with a 4.2 percentage point increase in the likelihood of reporting income near that kink – a 24% increase over the baseline likelihood of 18%. Under-reporting wages is significantly correlated with three kinks, but the effects there are smaller (about 0.4–1.2 percentage points). We also find that taxpayers who report wage income but do not have an associated Form W-2 are disproportionately likely to locate at the kinks examined in Table 5, with the strongest correlation occurring at the refund-maximizing kink.

The correlations for the self-employed, unsurprisingly, are even stronger. Reporting self-employment income increases the probabilities of locating near the refund-maximizing kink and the first EITC kink by 7.1 and 8.2 percentage points, respectively, relative to baseline likelihoods

Table 5: Regression analysis of taxpayers near kinks (2005-2014)

	Refund Max.	First EITC	Second EITC	CTC Ref.
Single	2.12	2.70	-0.06	1.17
	(0.10)	(0.10)	(0.09)	(0.10)
Self-employed	7.08	8.21	0.32	5.67
	(0.11)	(0.10)	(0.09)	(0.13)
Self-employed, no 1099-MISC	7.08	6.12	-1.37	3.93
	(0.12)	(0.11)	(0.13)	(0.15)
Wage earner, no W-2	2.31	1.76	0.41	1.32
	(0.21)	(0.22)	(0.20)	(0.20)
Over-report wages	3.86	1.99	2.37	4.24
	(0.14)	(0.15)	(0.13)	(0.16)
Under-report wages	1.01	1.16	0.04	0.38
	(0.10)	(0.10)	(0.10)	(0.13)
Paid preparer	-0.11	0.03	0.16	1.48
	(0.21)	(0.21)	(0.18)	(0.36)
Paper self-prepared	0.78	1.38	0.08	1.74
	(0.24)	(0.24)	(0.21)	(0.36)
Electronic self-prepared	1.42	1.60	0.26	2.79
	(0.21)	(0.21)	(0.18)	(0.30)
1 child	6.59	7.18	-0.12	—
	(0.09)	(0.08)	(0.08)	
2 children	3.21	3.73	-0.12	-4.93
	(0.09)	(0.09)	(0.08)	(0.09)
3+ children	2.39	2.76	0.85	-4.32
	(0.13)	(0.14)	(0.12)	(0.12)
Observations	1,642,808	1,666,799	1,638,825	1,107,776
Clusters (taxpayers)	809,400	802,318	834,653	585,175
Baseline probability of locating near the kink	19.26 (0.03)	19.30 (0.03)	14.59 (0.03)	17.98 (0.04)

The four columns of this table display the results of probit regressions where the dependent variable is an indicator for whether a taxpayer reports income near a given kink. The independent variables include those listed in the table as well as state and year fixed effects, and indicators for changing states, age (grouped into decade bins), positive and negative changes in the number of dependents claimed, getting married, and getting divorced. The coefficients are average marginal effects, presented as percentages, with standard errors in parentheses. The samples in each regression are limited to taxpayers with earned income within $3,500 (in 2014 dollars) of that column's kink. The first three columns are further limited to EITC-eligible taxpayers. Wage over-reporting and under-reporting variables are constructed by comparing wages as reported on Form 1040 and Form W-2. Errors are clustered at the taxpayer level. The final row contains the probability of locating near that column's kink, with the same sample restrictions as the probit regressions and with bootstrapped standard errors presented in parentheses. All figures reflect the authors' calculations using data from the IRS Compliance Data Warehouse.

strongest correlation at the refund-maximizing kink.

One way taxpayers may learn about kink incentives is through paid tax preparers. However, our analysis suggests usage of paid preparers is only meaningfully correlated with locating near

the CTC kink, with a relatively modest coefficient of 1.5 percentage points (an 8% increase over the baseline probability). Instead, the form of filing most associated with locating near kinks is electronic self-filing, suggesting that taxpayers may learn about kink incentives through the use of tax-filing software.

Family structure is also correlated with bunching behavior. In general, singles are more likely to bunch than married taxpayers. Those with one child are more likely to report incomes near the first EITC kink, the CTC kink, and the refund-maximizing kink. The second EITC kink, in contrast, attracts those with three or more children the most.

Whereas Table 5 presents a static snapshot of bunchers' characteristics, Table 6 explores two dynamic aspects of bunching: tracking the refund-maximizing kink and persistence at kinks over time. The columns on the left of Table 6 study the former, displaying the probability of reporting income within \$250 (in 2014 dollars) of the refund-maximizing kink in year t conditional on the taxpayer's distance from the refund-maximizing kink in year $t-1$. Importantly, the sample is limited here to taxpayers who saw the refund-maximizing kink move by at least \$2,000 (in 2014 dollars) between years $t-1$ and t. After experiencing such a large change in the refund-maximizing kink's location, roughly 23% of self-employed taxpayers who were near the refund-maximizing kink in year $t-1$ manage to locate near the kink again in year t. In contrast, self-employed taxpayers who were at least \$750 (in 2014 dollars) away from the refund-maximizing kink in year $t-1$ have between a 13% and 18% chance of locating near the refund-maximizing kink in year t.

The evidence here complements that of Figures 7 and 8, which show large shifts in the income distribution given abrupt changes in the location of the refund-maximizing kink. In the figures, it is unclear whether individual taxpayers move directly from one kink to another. The evidence we present in Table 6 indicates this is indeed the case.

Table 6 also highlights a recurring theme throughout the paper: the difference in bunching patterns between those reporting zero and non-zero self-employment income. The baseline probability of locating near the refund-maximizing kink in year t is much lower for wage earners – between 4% and 5% compared to 13% to 23% for the self-employed – because the propensity of wage earners to locate near any kink, including the refund-maximizing kink, is significantly lower. However, wage earners near the kink in year $t-1$ are about one half a percentage point more likely to report income near the refund-maximizing kink in year t than the other income groups analyzed here.

The columns on the right side of Table 6 test whether taxpayers persist at kinks over time. The last column displays average marginal effects, reported as percentage changes, from a probit regression where the dependent variable is an indicator of whether the taxpayer located in the same \$1,000-wide (in 2014 dollars) income bin in the current and previous year. The income bins range from \$5,000 to \$49,000 (in 2014 dollars) of earned income, and four indicator variables are also constructed for whether the income bin contains one of four kinks. These indicators are included as independent regressors, with additional regressors consisting of various household characteristics and interaction terms between the kink dummies and a self-employment dummy.

Table 6: Regression analysis of dynamic responses to kinks (2005-2014)

Follow the refund-maximizing kink?			Persist in income bin?	
Income in $t-1$ relative to kink	Wage earners	Self-empl.	Dependent variable: $\mathbf{1}\{\text{bin}_{t-1}=\text{bin}_t\}$	
$[-\$2{,}250, -\$1{,}750]$	4.27	13.62	Near ref. max. in $t-1$	0.12
	(0.08)	(0.27)		(0.12)
$[-\$1{,}750, -\$1{,}250]$	4.09	14.35	Near 1st EITC in $t-1$	0.67
	(0.08)	(0.26)		(0.12)
$[-\$1{,}250, -\$750]$	4.27	16.04	Near 2nd EITC in $t-1$	1.62
	(0.09)	(0.25)		(0.10)
$[-\$750, -\$250]$	4.60	19.18	Near CTC in $t-1$	1.16
	(0.09)	(0.22)		(0.12)
$[-\$250, +\$250]$	4.97	23.22	Self $\times$ ref. max.	4.42
	(0.09)	(0.18)		(0.20)
$[+\$250, +\$750]$	4.47	18.16	Self $\times$ 1st EITC	10.81
	(0.09)	(0.24)		(0.18)
$[+\$750, +\$1{,}250]$	4.29	16.93	Self $\times$ 2nd EITC	-0.15
	(0.09)	(0.27)		(0.21)
$[+\$1{,}250, +\$1{,}750]$	4.65	16.80	Self $\times$ CTC	1.88
	(0.09)	(0.30)		(0.20)
$[+\$1{,}750, +\$2{,}250]$	4.59	17.81	Self $\times$ no 1099-MISC	4.35
	(0.09)	(0.30)		(0.08)
			Wage $\times$ no W-2	6.61
				(0.14)
Observations	498,250	219,925	Observations	4,296,601
			Clusters (taxpayers)	1,101,184

The first three columns summarize the probability of reporting income within \$250 of the refund-maximizing kink in year t conditional on year $t-1$ income, measured relative to the refund-maximizing kink in year $t-1$. The sample is limited to EITC-eligible taxpayers whose refund-maximizing kink moved by at least \$2,000 between years $t-1$ and t, whose self-employment status remained constant in both years, and who were between 20 and 80 years of age in year t. The data reflect a 10 percent sample of all taxpayers between 2005 and 2014. The fourth and fifth columns display probit regression results where the dependent variable is an indicator for whether the taxpayer reported income in the same \$1,000-wide income bin in years $t-1$ and t. The coefficients are average marginal effects, presented in percentage terms. The data here reflect a two percent sample of EITC-eligible taxpayers with income between \$5,000 and \$49,000, aged 20 to 80, from 2005 to 2014. Unreported independent variables include a set of indicator variables for filing status, number of dependents, change in filing status, change in number of dependents, self-employment status, preparer type, changes in state of residence, and the sign of the conflict in reported wages on Form W-2 and Form 1040 (if applicable). Year, age, and state fixed effects are also included in the regression. Errors are clustered at the taxpayer level. All dollar amounts in the table and in this note are given in 2014 dollars. All figures reflect the authors' calculations using data from the IRS Compliance Data Warehouse.

The regression results indicate persistence is higher near the first EITC, second EITC, CTC refundability, and refund-maximizing kinks than income bins not containing any of these kinks. Near the kinks, the self-employed persist to a greater degree than wage earners, except at the second EITC kink. Self-employed individuals near the first EITC kink are the most persistent, followed by the those near the refund-maximizing kink. The self-employed at the CTC kink persist at a relatively weaker level, and the weakest overall persistence occurs at the second EITC kink. Those

without third-party reports documenting their earnings – i.e., self-employed taxpayers with no 1099-MISC income and wage earners with no W-2 income – are much more likely to persist than those with third-party reporting.

IV Discussion

Our investigation of taxpayer bunching around kink points in the United States income tax schedule has revealed new details about the way (some) taxpayers respond to the tax code. Having examined federal and state statutory kinks, kinks created by tax credits, and kinks created by phase-outs of deductions and exemptions, we have seen that most kinks do not generate statistically discernible bunching responses. In particular, all non-convex kinks fail to induce an observable response, and we do not detect any responsiveness among high-income taxpayers at the federal or state income tax kinks we study, even when using the universe of tax returns.

However, we do find economically meaningful bunching at several large, convex kinks in the tax schedule, with statistically significant bunching occurring at a total of seven kinks. Consistent with Saez (2010) and Chetty, Friedman and Saez (2013), the first EITC kink elicits the strongest response. Bunching there occurs during all years of our sample, 1996 to 2014, and increases in intensity over the course of the sample. We also find evidence of new bunching responses, most notably at the CTC refundability plateau and the second EITC kink. Bunching there does not occur at the beginning of our sample but emerges during the mid 2000s. In addition, we find small but visually compelling, statistically significant responses at the largest statutory kink, the beginning of the statutory schedule, the third statutory kink, and the temporary Making Work Pay Tax Credit kink.

Bunching is growing more common, with the total number of bunchers at the kinks we study growing from roughly 104,000 in 1996 to about 834,000 in 2014. This represents a 700% increase, despite the tax filing population's modest increase of 20%, suggesting that taxpayers may be actively learning about the incentives embedded in the tax code. However, our investigations into the mechanism behind this learning process yield little fruit. Neither the use of paid preparers nor tax preparation software are strongly associated with bunching at any of the kinks we study. Regardless of how taxpayers learn to bunch, however, there is a significant tendency to persist at kinks once bunching is initiated. Relative to placebo income bins, bins containing kinks are more likely to see taxpayers remain inside the bin from year to year.

Though the self-employed represent only 12% of taxpayers in our sample, we estimate they comprise 89% of bunchers. Because self-employment income is known to suffer from inaccurate reporting, this raises the concern that bunching is driven by strategic misreporting of income. This explanation is consistent with two facts we have documented. First, bunching tends to occur at kinks that maximize the tax credits taxpayers receive, but not at other large kinks in the tax schedule. Second, taxpayers whose self-employment income is reported by third parties on

Form 1099-MISC are less likely to bunch. As further evidence in support of this hypothesis, Kuka (2013) shows that expansions of the EITC in the mid 1990s led to an extensive margin response in taxpayer-reported self-employment income that is not present in the Census' Current Population Survey. However, it remains possible that a large share of self-employment bunching is due to legitimate control over labor income. Further research is needed to definitively answer this question.

For the case of wage earners, misreporting can be detected directly, and here the evidence is clear: bunching patterns do not manifest in employer-reported wages. Further, mismatches between taxpayer- and employer-reported wages are a strong predictor of locating near a kink. Typically, a buncher without self-employment income reports *more* wage income than their employer reports on Form W-2. The implication is that wage-earner bunching patterns merely reflect reporting decisions, not distortions in real economic activity. Whether this strategic misreporting behavior is successful, however, remains unclear, as we do not observe post-audit outcomes.

Our findings raise theoretical concerns about the conventional mapping from bunching patterns to elasticities. At first blush, standard labor supply models are entirely consistent with bunching behavior. In such models, taxpayers adjust labor hours continuously, and kinks create discontinuities in marginal incentives, causing taxpayers with a range of preferences to optimally earn income at the kink. Building on this insight, Saez (2010) derives a one-to-one mapping between the strength of the bunching pattern and the elasticity of taxable income with respect to the marginal tax rate. In principle, elasticities calculated in this way can capture the sensitivity of labor supply, deduction, and evasion decisions to marginal tax rates.

However, the bunching patterns we discuss in this paper are often inconsistent with the basic assumptions underlying standard models. At times, taxpayer bunching seems oddly disproportionate to the kink size (e.g., in Figure 4). In other cases, small changes to the effective tax schedule cause large shifts in the distribution of income (e.g., in Figures 7 and 8). We have argued that this behavior is due to taxpayers reporting incomes at the refund-maximizing kink – that is, the kink that maximizes tax credits net of taxes owed. This is wholly inconsistent with models that translate bunching patterns into elasticities, as refund-maximizing taxpayers are insensitive to changes in marginal tax rates per se. Instead, these taxpayers' reported incomes respond if and only if the refund-maximizing kink changes location. For this reason, we have omitted elasticity estimates from the paper. Exploring the theoretical implications of such target-seeking behavior remains an interesting avenue for future research.

References

Allingham, Michael G., and Agnar Sandmo. 1972. "Income Tax Evasion: A Theoretical Analysis." *Journal of Public Economics*, 1(3-4): 323–338.

Bastani, Spencer, and Håkan Selin. 2014. "Bunching and non-bunching at kink points of the Swedish tax schedule." *Journal of Public Economics*, 109(C): 36–49.

Blomquist, Sören, and Whitney Newey. 2002. "Nonparametric estimation with nonlinear budget sets." *Econometrica*, 70(6): 2455–2480.

Burtless, Gary, and Jerry A. Hausman. 1978. "The effect of taxation on labor supply: Evaluating the Gary negative income tax experiment." *The Journal of Political Economy*, 1103–1130.

Burtless, Gary, and Robert Moffitt. 1984. "The effect of Social Security benefits on the labor supply of the aged." *Retirement and Economic Behavior*, 135–171.

Chetty, Raj. 2009. "Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance." *American Economic Journal: Economic Policy*, 1(2): 31–52.

Chetty, Raj. 2012. "Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply." *Econometrica*, 80(3): 969–1018.

Chetty, Raj, and Emmanuel Saez. 2013. "Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipients." *American Economic Journal: Applied Economics*, 5(1): 1–31.

Chetty, Raj, John N. Friedman, and Emmanuel Saez. 2013. "Using Differences in Knowledge across Neighborhoods to Uncover the Impacts of the EITC on Earnings." *American Economic Review*, 103(7): 2683–2721.

Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri. 2011. "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records." *The Quarterly Journal of Economics*, 126(2): 749–804.

Crandall-Hollick, Margot. 2013. "The Child Tax Credit: Current Law and Legislative History." *Congressional Research Service Report*, 41873.

Devereux, Michael P., Li Liu, and Simon Loretz. 2014. "The elasticity of corporate taxable income: New evidence from UK tax records." *American Economic Journal: Economic Policy*, 6(2): 19–53.

Einav, Liran, Amy Finkelstein, and Paul Schrimpf. 2016. "Bunching at the Kink: Implications for Spending Responses to Health Insurance Contracts." NBER Working Paper 22369.

Eissa, Nada, and Hilary Hoynes. 2011. "Redistribution and tax expenditures: The Earned Income Tax Credit." *National Tax Journal*, 64(2): 689–730.

Feenberg, Daniel, and Elisabeth Coutts. 1993. "The TAXSIM Model." *Journal of Policy Analysis and Management*, 12(1): 189–194.

Feldman, Naomi E., Peter Katuscak, and Laura Kawano. 2015. "Taxpayer Confusion: Evidence from the Child Tax Credit." *Available at SSRN 2296542*.

Foundation, Kaiser Family. 2015. "Annual Updates on Eligibility Rules, Enrollment and Renewal Procedures, and Cost-Sharing Practices in Medicaid and CHIP." kff.org/medicaid/report/annual-updates-on-eligibility-rules-enrollment-and/ Accessed: 2015-09-04.

Friedberg, Leora. 2000. "The labor supply effects of the Social Security earnings test." *Review of Economics and Statistics*, 82(1): 48–63.

Gelber, Alexander M., Damon Jones, and Daniel W. Sacks. 2013. "Earnings Adjustment Frictions: Evidence from the Social Security Earnings Test." NBER Working Paper 19491.

Hargarden, Enda Patrick. 2015. "Taxpayer Responses over the Cycle: Evidence from Irish Notches."

Hoxby, Caroline M., and George B. Bulman. 2016. "The effects of the tax deduction for post-secondary tuition: Implications for structuring tax-based aid." *Economics of Education Review*, 51: 23–60.

Internal Revenue Service. 2014. "Compliance Estimates for the Earned Income Tax Credit Claimed on 2006 - 2008 Returns, Publication 5162." 66766H.

Internal Revenue Service. 2016. "Tax Gap Estimates for Tax Years 2008-2010."

Kleven, Henrik J. 2016. "Bunching." *Annual Review of Economics*, 8(1).

Kleven, Henrik Jacobsen, Martin B. Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez. 2011. "Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark." *Econometrica*, 79(3): 651–692.

Kleven, Henrik J., and Mazhar Waseem. 2013. "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan." *The Quarterly Journal of Economics*, 128(2): 669–723.

Kuka, Elira. 2013. "EITC and the Self-employed: Real or Reporting Effects?" *Public Finance Review*, 1091142113496130.

Le Maire, Daniel, and Bertel Schjerning. 2013. "Tax bunching, income shifting and self-employment." *Journal of Public Economics*, 107(C): 1–18.

Nichols, Austin, and Jesse Rothstein. 2015. "The Earned Income Tax Credit." NBER Working Paper 21211.

Ramnath, Shanthi. 2013. "Taxpayers' Responses to Tax-based Incentives for Retirement Savings: Evidence from the Saver's Credit Notch." *Journal of Public Economics*, 101(C): 77–93.

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

Slemrod, Joel. 1985. "An empirical test for tax evasion." *The Review of Economics and Statistics*, 232–238.

Slemrod, Joel. 2007. "Cheating ourselves: The economics of tax evasion." *The Journal of Economic Perspectives*, 21(1): 25–48.

Appendix A: Bunching at the Zeroth and Third Statutory Kinks

Here we present figures showing bunching at the zeroth and third statutory kinks. The latter sees bunching only by married taxpayers filing separately, and only in a handful of years in our sample. Figure A1 depicts the full population of tax returns for this group, with data pooled from 1996 to 2014.

Figure A1: Bunching at the third statutory kink: Married-filing-separately (1996-2014)

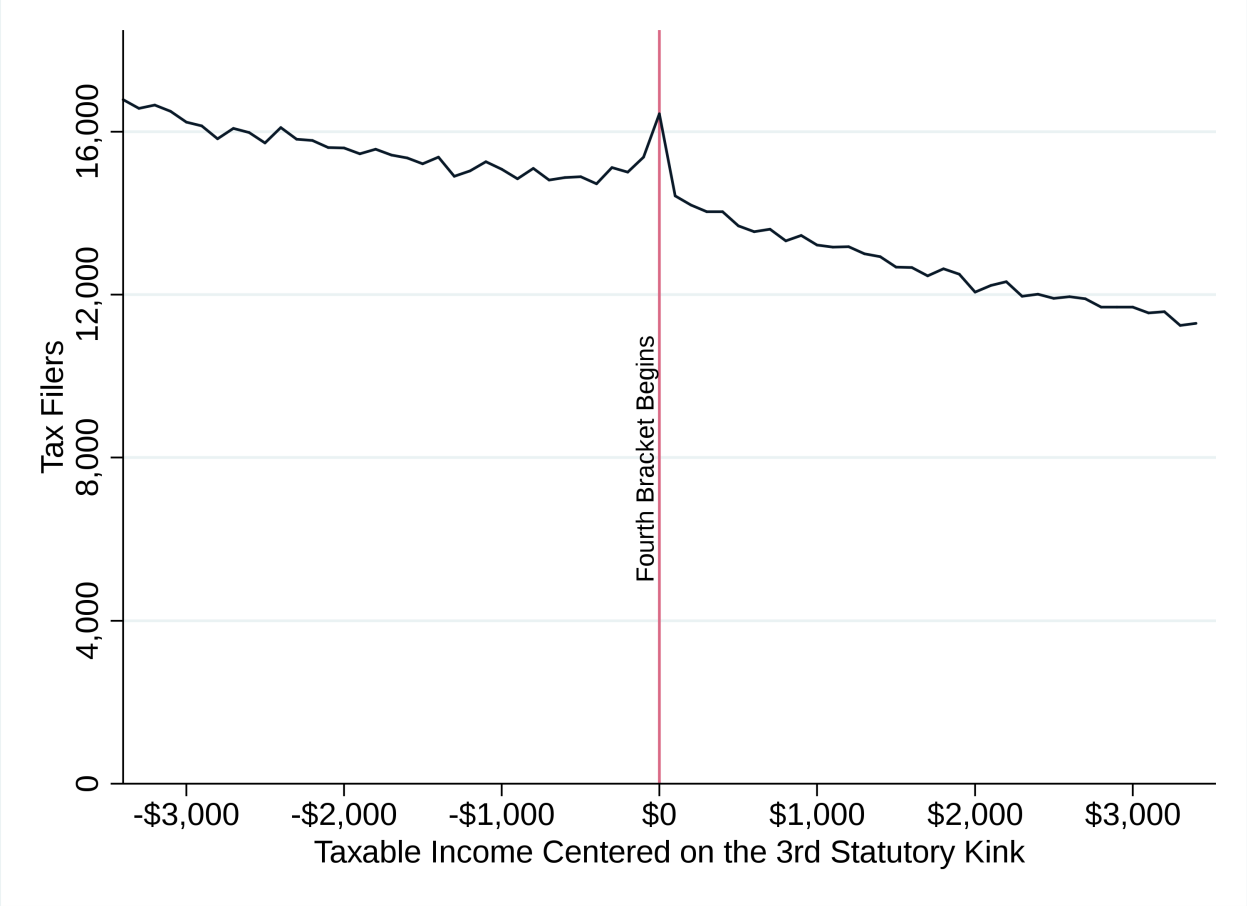

The distribution of taxable income is displayed for the full population of married couples filing separately from 1996-2014. Income is rounded into $100 bins, and the distribution is centered on the third statutory kink. All dollar values are inflation adjusted to 2014 levels. This figure was created by the authors using data from the IRS Compliance Data Warehouse.

As discussed in Section I.B, 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 generally do not face a kink here, due to complications with the Child Tax Credit. For this reason, we only analyze childless, self-employed taxpayers at this kink. The income distribution centered around this kink is displayed separately

for single filers and married filing jointly filers in Figures A2 and A3. These figures are the product of pooling data from 2002 to 2014.

Figure A2: Bunching at the zeroth statutory kink: Single, self-employed, childless taxpayers (2002-2014)

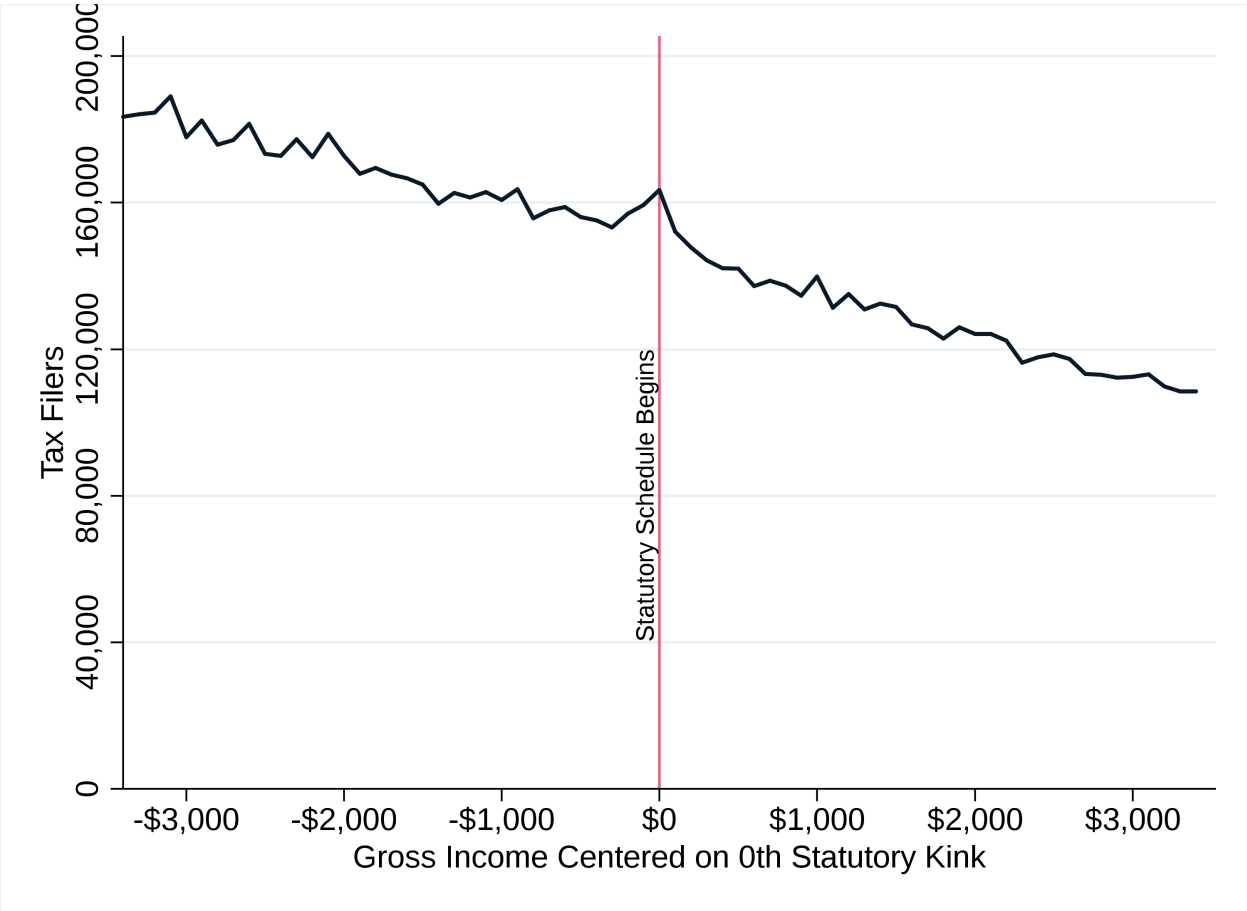

The distribution of gross income is displayed for singles filing from 2002-2014 with self-employment income and no dependents. Income is rounded into $100 bins, and the distribution is centered on the zeroth statutory kink. All dollar values are inflation adjusted to 2014 levels. This figure was created by the authors using data from the IRS Compliance Data Warehouse and has been weighted to represent the full U.S. population.

Figure A3: Bunching at the zeroth statutory kink: Married-filing-jointly, self-employed, childless taxpayers (2002-2014)

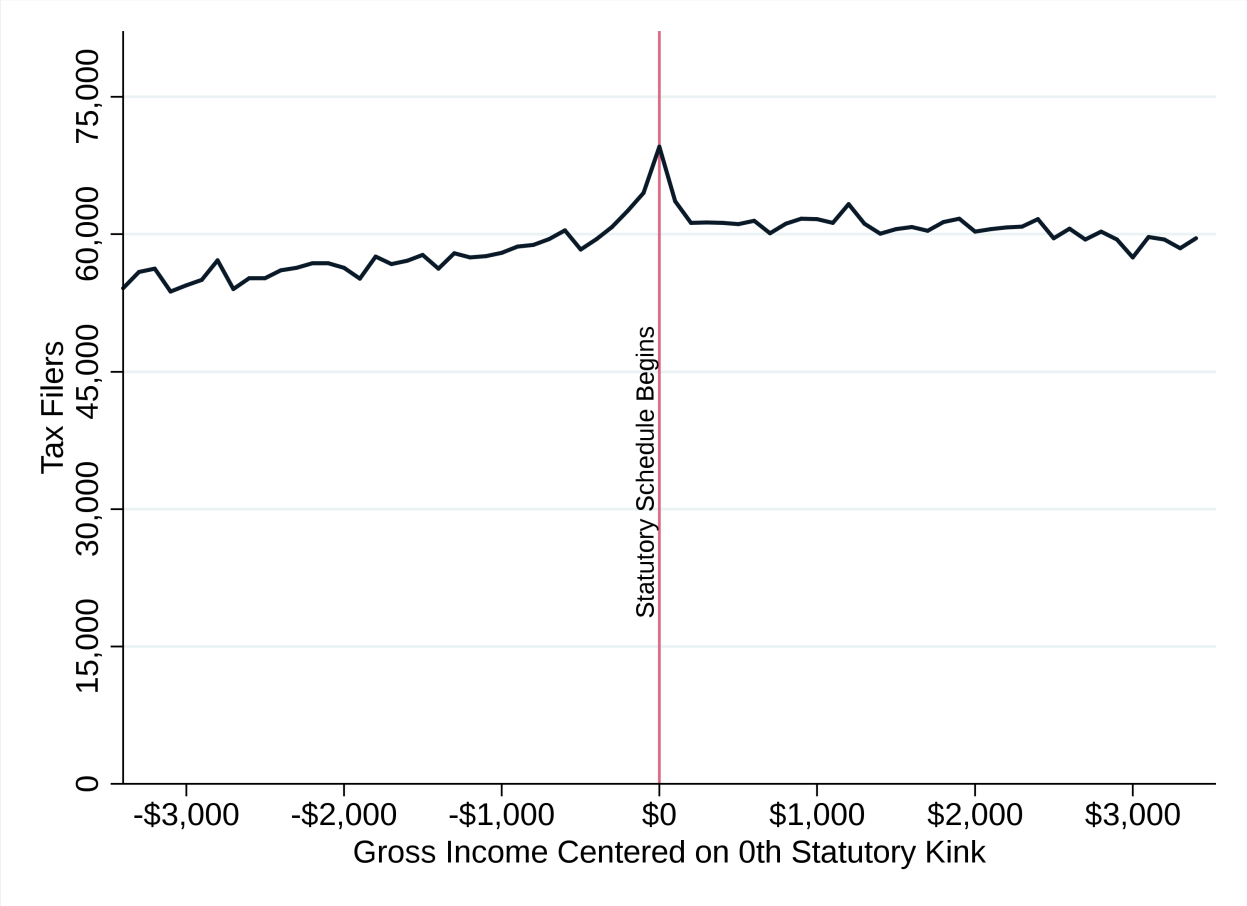

The distribution of gross income is displayed for married couples filing jointly from 2002-2014 with self-employment income and no dependents. Income is rounded into $100 bins, and the distribution is centered on the zeroth statutory kink. All dollar values are inflation adjusted to 2014 levels. This figure was created by the authors using data from the IRS Compliance Data Warehouse and has been weighted to represent the full U.S. population.

FOR ONLINE PUBLICATION

Appendix B: Robustness Checks

Here we test our estimation technique for sensitivity to parameter choice. The three key parameters are binwidth and the sizes of the bunching window and bunching region. In Section II.A, we label these parameters δ, W, and R, respectively. Binwidth simply 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.

As an example, our default parameter values are $\delta = \$100$, $W = 10$, and $R = 35$. This implies a bunching window of $W \cdot \delta = \$1,000$ and a bunching region of $R \cdot \delta = \$3,500$ around the kink. In other words, we assume the distribution of income within \$1,000 of the kink is affected by bunching, but that outside this threshold the distribution is unaffected by bunching. Moreover, we use the observed distribution of income between \$1,000 and \$3,500 away from the kink to estimate the counterfactual distribution of income if the kink did not exist.

Tables A1 and A2 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 Section II.A. By presenting 2003 estimates here, we avoid this issue and thus are able to test a wide range of parameter choices.

The results indicate 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 53.6%. Binwidth choices of \$50 or \$250 lead to estimates of 53.6% and 52.0%, respectively, with small accompanying standard errors. Changing the bunching region by \$500 in either direction also has small effects, with alternative estimates of 52.2% and 55.0%, again with small standard errors. The one parameter that has a substantive effect on the estimates is the choice of bunching window. In particular, expanding the bunching window to \$1,500 causes the bunching coefficient to fall to 49.6%. This is unsurprising, as visual inspection of panel (a) of Figure A4 makes clear that far fewer taxpayers between \$1,000 and \$1,500 away from the kink should be classified as bunchers. More surprising is the fact that reducing the bunching window from \$1,000 to \$500 only increases the bunching coefficient by 1.5 percentage points. This seems wrong, as visual inspection clearly shows that the fraction of bunchers within \$500 of the kink is much higher than the fraction within \$1,000 of the kink. However, the reason the bunching coefficient does not rise more is that \$500 is not a sensible choice for a bunching window. The distribution of income just outside of \$500 away from the kink is substantially distorted due to bunching, leading to artificially high estimates for the counterfactual distributions of income near the kink.

Other groups show similar patterns. Our preferred bunching coefficient estimate for single, self-employed taxpayers with two children is 43.4%. Except for the choice of bunching window, all

Table A1: Bunching coefficients calculated at the first EITC kink in 2003

	Bunching coefficient	Sample size	Bin-width	Bunching window	Bunching region
Single, self-employed, one child	53.6% (1.3%)	42,900	$100	$1,000	$3,500
	55.1% (1.9%)	42,900	$100	$500	$3,500
	49.6% (1.7%)	42,900	$100	$1,500	$3,500
	52.2% (1.6%)	40,000	$100	$1,000	$3,000
	55.0% (1.3%)	45,600	$100	$1,000	$4,000
	53.6% (1.3%)	42,900	$50	$1,000	$3,500
	52.0% (1.4%)	43,400	$250	$1,000	$3,500
Single, self-employed, two children	43.4% (1.5%)	37,000	$100	$1,000	$3,500
	48.2% (1.6%)	37,000	$100	$500	$3,500
	43.3% (1.5%)	37,000	$100	$1,500	$3,500
	41.6% (1.8%)	34,800	$100	$1,000	$3,000
	46.3% (1.5%)	39,200	$100	$1,000	$4,000
	43.4% (1.5%)	36,800	$50	$1,000	$3,500
	44.1% (1.3%)	37,300	$250	$1,000	$3,500

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 three 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 IRS Compliance Data Warehouse.

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

	Bunching coefficient	Sample size	Bin-width	Bunching window	Bunching region
Married, self-employed, one child	37.4% (1.7%)	11,400	$100	$1,000	$3,500
	43.3% (1.7%)	11,400	$100	$500	$3,500
	31.7% (2.3%)	11,400	$100	$1,500	$3,500
	37.7% (2.1%)	10,300	$100	$1,000	$3,000
	38.0% (1.5%)	12,500	$100	$1,000	$4,000
	37.4% (1.7%)	11,400	$50	$1,000	$3,500
	35.7% (1.6%)	11,700	$250	$1,000	$3,500
Married, self-employed, two children	30.9% (1.4%)	44,100	$100	$1,000	$3,500
	36.2% (1.4%)	44,100	$100	$500	$3,500
	27.9% (1.8%)	44,100	$100	$1,500	$3,500
	29.0% (1.6%)	40,000	$100	$1,000	$3,000
	32.8% (1.4%)	47,800	$100	$1,000	$4,000
	30.9% (1.6%)	43,900	$50	$1,000	$3,500
	31.0% (1.1%)	44,600	$250	$1,000	$3,500

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 three 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 IRS Compliance Data Warehouse.

Figure A4: Income distribution of self-employed taxpayers near the first EITC kink in 2003

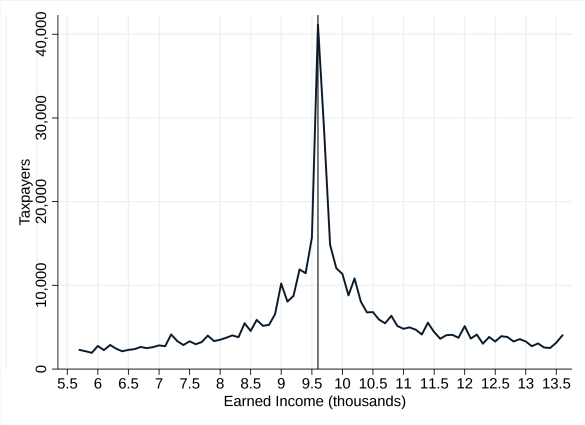

(a) Single, one child

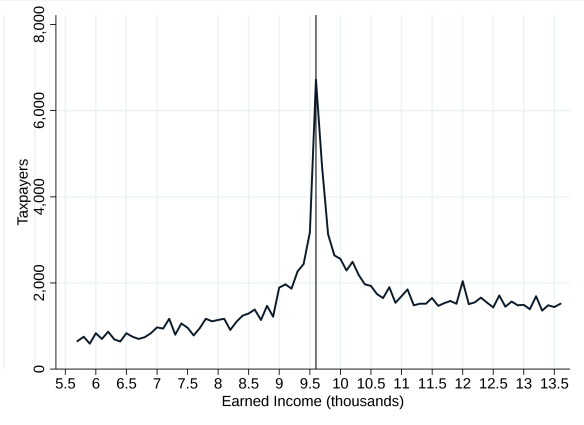

(b) Married filing jointly, one child

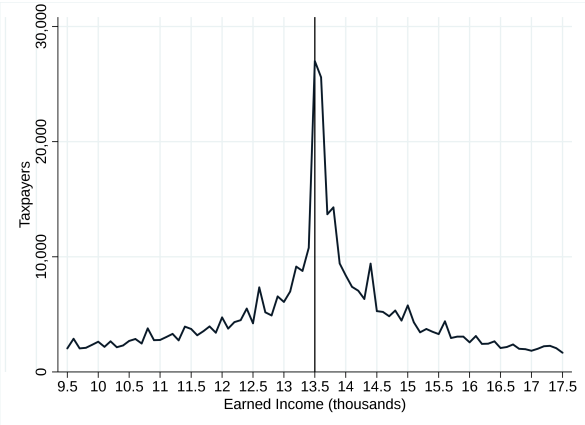

(c) Single, two children

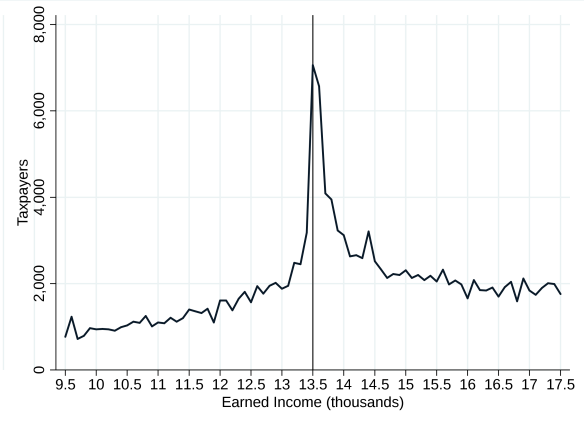

(d) Married filing jointly, two children

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 IRS Compliance Data Warehouse and have been weighted to represent the full U.S. population.

other permutations leave the bunching coefficient in the range $[41.6\%, 46.3\%]$. For these taxpayers a smaller bunching window of \$500 increases the estimate to 48.2%, and a larger bunching window has little effect. For married, self-employed taxpayers, our preferred bunching coefficient estimates are 37.4% and 30.9%, respectively, for those with one or two children. Except for the choice of bunching window, alternative estimates for these parameters lie in the ranges $[35.7\%, 38.0\%]$ and $[29.0\%, 32.8\%]$, respectively. For all cases, the bunching window could arguably be expanded from \$1,000 to \$1,500. This would generally decrease our bunching coefficients, but it would change the interpretation of the coefficient (from percentage within \$1,000 who are bunching to percentage within \$1,500 who are bunching) such that the drop in bunching coefficients would not imply a lesser degree of bunching.

FOR ONLINE PUBLICATION

Appendix C: Bunching Analysis of Medicaid, SNAP, and Federal Disability Insurance

In this appendix we analyze whether federal income tax data provide evidence that individuals adjust their incomes to remain eligible for Medicaid, Supplemental Nutrition Assistance Program (SNAP), or federal disability benefits. These are large, economically important programs that serve millions of people. In 2011, there were approximately 57 million Medicaid participants, 45 million SNAP beneficiaries, and 8.5 million workers receiving federal disability benefits. We analyze these programs as a robustness check to ensure perceived bunching at kink points in the tax schedule was not caused by other incentives. We describe each program in broad strokes, with an emphasis on program-specific income definitions and eligibility thresholds. We then discuss our findings and conclude with a short discussion of the strengths and weaknesses of our approach. In short, we see no evidence of bunching associated with any of these programs, but this may be due to the limitations of the tax data in the context of these programs.

C.1 A Brief Description of Each Program

Medicaid provides health insurance at subsidized rates to low-income individuals, primarily parents, pregnant mothers, and children. For many, the program is free. Income eligibility criteria are a function of the federal poverty line, which is an increasing function of the number of adults and children in the household. Medicaid is administered at the state level, and the definitions of qualifying income and eligibility thresholds (as a percentage of the federal poverty line) vary by state. Importantly, during the period we analyze (2002 to 2011) Medicaid introduces a "notch" in the budget set, as earning income above the threshold results in a complete loss of benefits. This is in contrast to having benefits phase-out, which would produce a kink. Given that the threshold is a notch and that Medicaid is a large benefit, we expect substantial responsiveness to this threshold.

SNAP provides funds to low-income individuals that can only be spent on certain types of food at participating retailers. Similar to Medicaid, it is a function of the federal poverty line, with larger households receiving greater benefits. Unlike Medicaid, the thresholds and income definitions do not vary across states, with the exceptions of Alaska and Hawai'i. Eligibility is limited to those with less than 130% of the federal poverty line in gross monthly income and less than 100% of the poverty line in net monthly income. Benefits begin phasing out with the first dollar of earned income. Individuals receive the difference between the maximum allotment per household (around $650 for a family of 4) and 30% of their monthly income. This produces a relatively modest non-convex kink at the end of the phase-out region.[31]

Federal disability insurance (DI) is administered by the Social Security Administration, providing monthly payments to individuals with a disability. Individuals must apply for DI and once approved are unable to earn income above a certain threshold without triggering a review of their

[31] See www.fns.usda.gov/snap/eligibility for a thorough description of SNAP, including eligibility criteria and definitions of income.

claim or outright termination of benefits. Thus, the threshold (called the Substantial Gainful Activity threshold) is a notch. The Substantial Gainful Activity threshold ranged from $500 each month in 1996 to $1,000 each month in 2011. The size of the monthly benefit for DI recipients is a function of prior earnings and can range from a few hundred dollars to a few thousand dollars per month. As a result, we believe the potential for bunching to the left of this notch is substantial.[32]

C.2 Qualifying Income and Kink Construction

SNAP and DI income eligibility thresholds are measured monthly, and in some states Medicaid income is measured monthly as well. Given that tax data record income on an annual basis, we construct the applicable kinks or notches by simply multiplying the monthly thresholds by twelve.[33]

All three programs define income in an analogous manner to "earned income" qualifying for the EITC, but can include income items that are not recorded by the tax system, such as child support payments, housing subsidies, or Supplemental Security Income. In addition, various deductions are allowed and SNAP beneficiaries must satisfy the gross and net income tests described above. Medicaid income varies by state, but we use the same income definition for all states and all years.

We use the EITC sample as the basis for studying all three programs. These data are drawn with the following restrictions: all observations are from the seven states with no income taxes, all filed federal income tax returns, and all were between 25 and 65 years of age. We analyze Medicaid bunching from 2002 to 2011, SNAP bunching from 1996 to 2011, and DI bunching from 1999 to 2011.

C.3 Bunching Results for Medicaid, SNAP, and Federal Disability Insurance

We find an absence of bunching at each eligibility threshold in every state and every year. This is surprising given the value of these programs to participants – especially in the case of Medicaid and DI. The lack of bunching suggests either individuals are not adjusting their income in response to these programs, or our estimates are biased by substantial measurement error. Individuals might not respond if they have imperfect knowledge of income eligibility criteria, or because adjusting income is costly.

On the other hand, individuals may be responding to these programs in ways that are undetectable with tax data. Federal income tax data do not contain all income or deduction items that comprise qualifying income for these programs. This is particularly relevant for the analysis of Medicaid, where the income definition varies across states and potentially over time. A second problem is that tax data are recorded annually, while eligibility for two (and in some states all three) programs are evaluated on a monthly basis. Thus, the only bunchers we can identify in our data for SNAP and DI are those that bunch in *every* month.

[32] See www.ssa.gov/disability for a thorough description of federal disability insurance.

[33] We retrieved annual state specific income eligibility thresholds for Medicaid from Foundation (2015).

FOR ONLINE PUBLICATION

Appendix D: 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 kinks created by phase-outs of personal exemptions, itemized deductions, and the American Opportunity Tax Credit.

D.1 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 AGI, generating notches in the budget constraint. Personal exemptions are reduced by 2% for each $2,500 of income exceeding the phase-out threshold until exemptions are exhausted. The beginning (and end) 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.[34]

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. However, 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% 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 $1,000 above the Pease threshold, qualified itemized deductions are reduced by 3%, meaning the individual has 30 additional dollars of taxable income. If the taxpayer faces an initial marginal tax rate of 31%, she would see her marginal rate increase by around 1 percentage point ($\$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. In the presence of moderate optimization frictions, these kinks are unlikely to induce a behavioral response.

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

[34]For all filers, the end of the phase-out region is $122,500 above the beginning.

of the discontinuity relative to a kink. For example, the "size" of a kink is calculated by dividing the difference between the net-of-tax rates (one minus the marginal tax rate) on either side of a kink, and dividing by the net-of-tax rate to the right of the kink. Calculating the size of a notch requires an assumption about the size of a marginal response by the taxpayer: do taxpayers adjust their income in $1, $50, or $1,000 increments?

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% 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.[35]

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.

D.2 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% 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 household with two students 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 households 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 cannot claim the credit and therefore we do not observe education expenses for them. For most households, in most years, there is no bunching response at the convex kink. For the few household 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.

[35]The formula for kink size, given an income increment of $X \leq 2500$, is $((79.36/X) \times 100\%)$.

Thus it provides an opportunity to test the hypothesis that high-income taxpayers would bunch if they faced incentives similar to the EITC and CTC. However, an important difference between the AOTC and the EITC and CTC is that the AOTC does not mark the refund-maximizing kink for any taxpayers in any year.

D.3 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).[36]

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.

[36]We caution, though, that New Jersey's definition of taxable income is different than the federal government's

FOR ONLINE PUBLICATION

Appendix E: Geographic Variation in Bunching

The body of this paper documents the evolution of bunching at various kinks over time. The trend – as displayed in Figure 5 – is that bunching is generally increasing over time. This appendix provides evidence of geographic variation in bunching at the refund-maximizing kink. It is a complement to the geographic variation in EITC maximization documented by Chetty, Friedman and Saez (2013).

Figure A5 displays bunching intensity by state, pooled for years 1996–2014, for EITC-eligible taxpayers with children and self-employment income. Here bunching intensity is the percentage of these taxpayers that report earned income within $500 of the refund-maximizing kink. The shades of color in the figure are deciles calculated by pooling all state-year combinations. These decile break-points are carried over to Figures A6 through A9, which display bunching intensity by state for 1996, 2003, 2009, and 2014, respectively.

In 1996 (Figure A6), bunching was weak relative to the rest of the sampled period, with only Florida and New York having greater than 9% of taxpayers near the refund-maximizing kink. By 2003 (Figure A7) more taxpayers were bunching in all states outside of the Great Plains. Bunching was particularly intense bunching in the southeast (Florida exceeded 20%), Texas (15%), and New York (18%). Bunching generally intensified for states outside of the southeast and New York between 2003 and 2009, the last year Chetty, Friedman and Saez (2013) cover.

In comparison with 1996, bunching in 2014 by EITC-eligible taxpayers with self-employment income and children had spread and intensified dramatically. Thirty-two states had greater than 10% of these taxpayers locate near the refund-maximizing kink, twenty-one exceeded 15% and five southeastern states exceeded 25%. Mississippi had the most intense bunching, with around 29% of these taxpayers locating near the refund-maximizing kink, compared to 6% for Mississippi in 1996 or 5% for South Dakota in 2014. These geographic and intertemporal patterns support earlier work suggesting bunching has a strong geographic component (Chetty, Friedman and Saez, 2013).

Figure A5: Bunching intensity at the refund-maximizing kink: 1996-2014

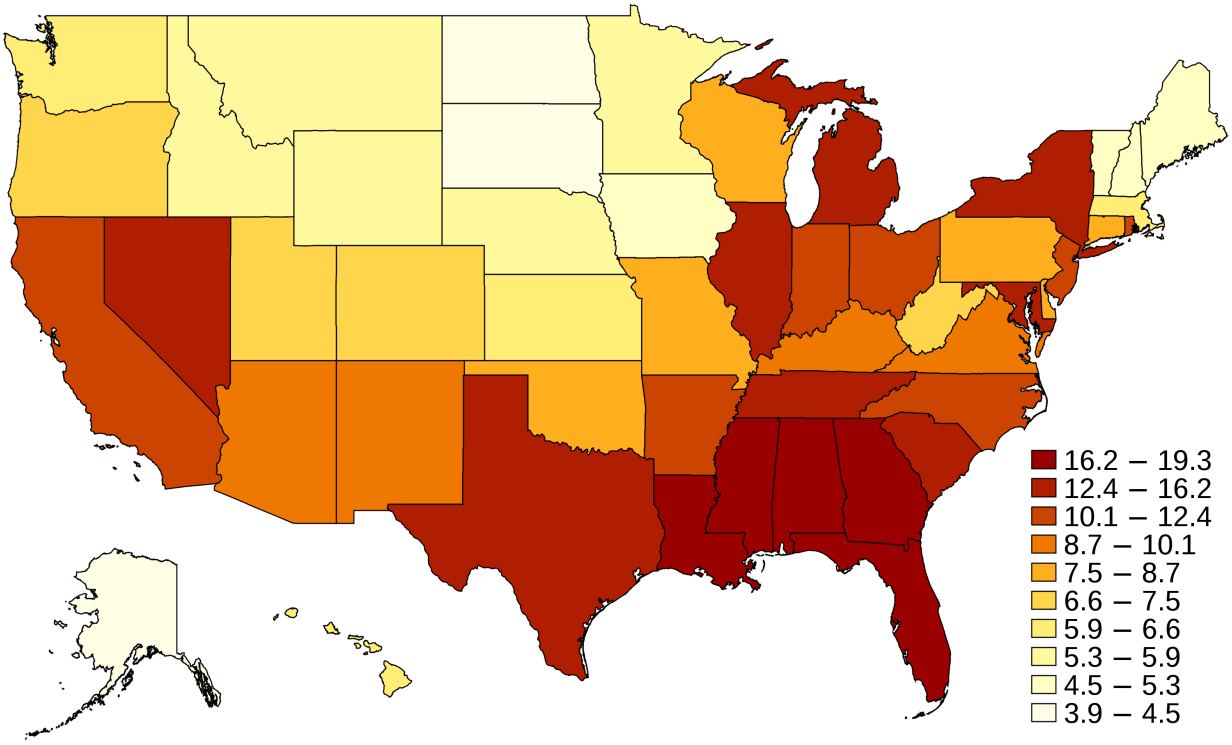

The sample is limited to taxpayers with children and self-employment income. Colors represent deciles of bunching intensities calculated by pooling all states and years. Intensity is proxied by the percentage of all EITC-eligible taxpayers that report earned income within $500 of the refund-maximizing kink in a given state. Dollar values are adjusted to 2014 levels. The District of Columbia is excluded. This figure was created by the authors using data from the IRS Compliance Data Warehouse.

Figure A6: Bunching intensity at the refund-maximizing kink: 1996

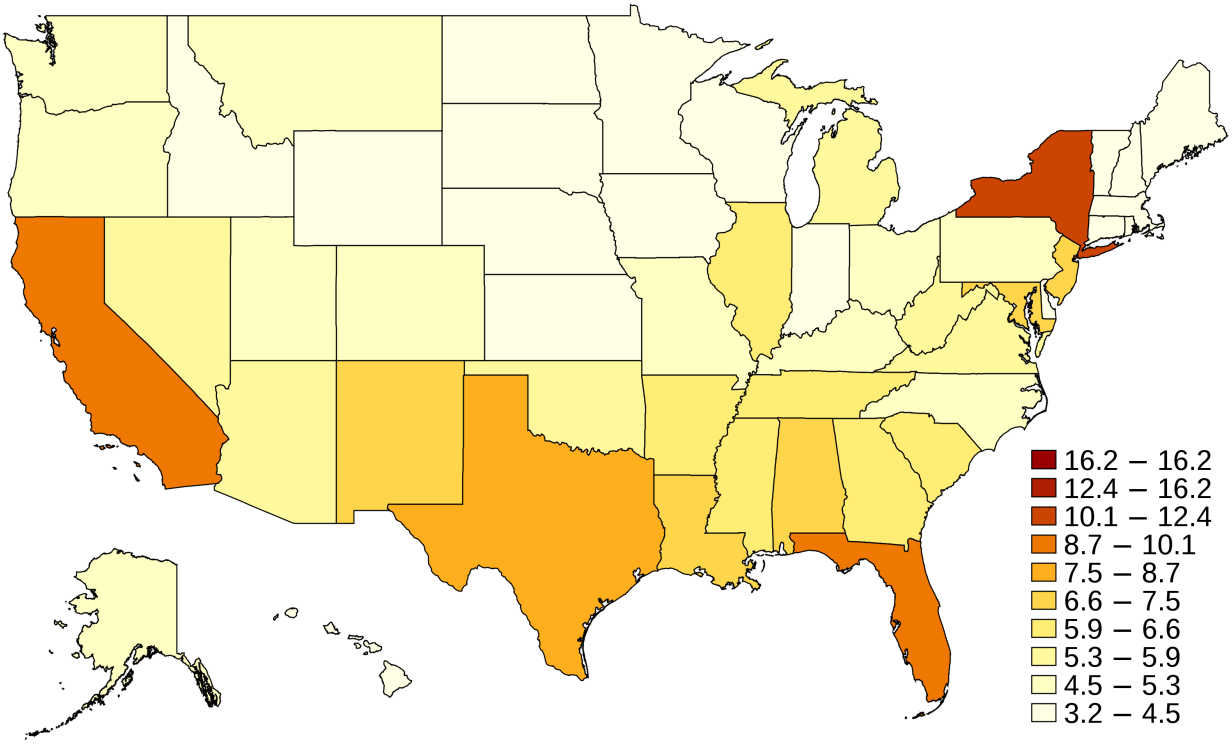

The sample is limited to taxpayers with children and self-employment income. Colors represent deciles of bunching intensities calculated by pooling all states and years. Intensity is proxied by the percentage of all EITC-eligible taxpayers that report earned income within $500 of the refund-maximizing kink in a given state. Dollar values are adjusted to 2014 levels. The District of Columbia is excluded. This figure was created by the authors using data from the IRS Compliance Data Warehouse.

Figure A7: Bunching intensity at the refund-maximizing kink: 2003

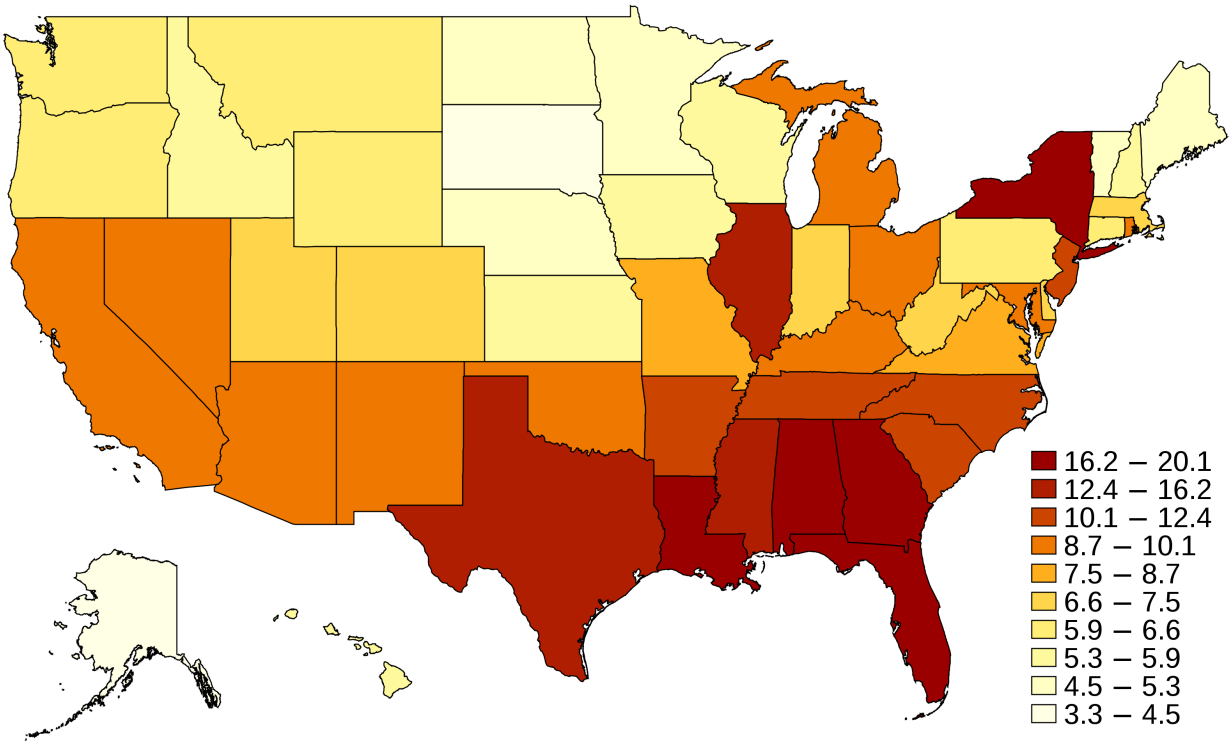

The sample is limited to taxpayers with children and self-employment income. Colors represent deciles of bunching intensities calculated by pooling all states and years. Intensity is proxied by the percentage of all EITC-eligible taxpayers that report earned income within $500 of the refund-maximizing kink in a given state. Dollar values are adjusted to 2014 levels. The District of Columbia is excluded. This figure was created by the authors using data from the IRS Compliance Data Warehouse.

Figure A8: Bunching intensity at the refund-maximizing kink: 2009

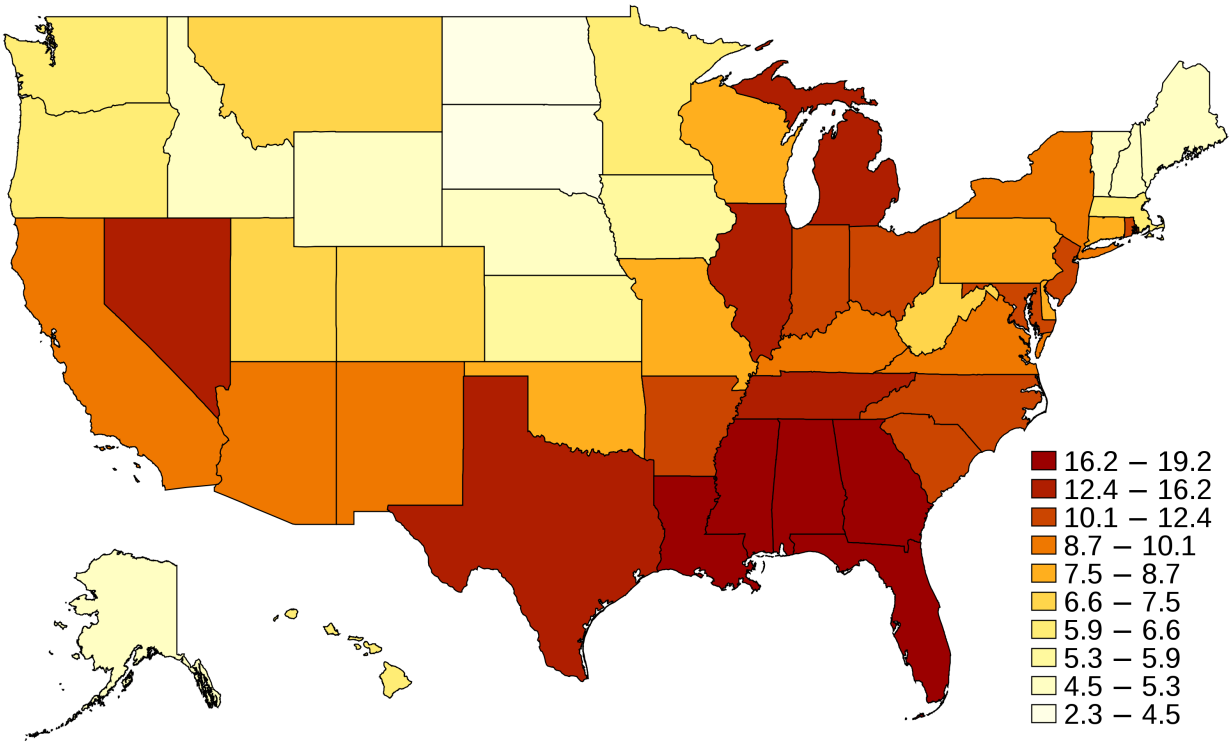

The sample is limited to taxpayers with children and self-employment income. Colors represent deciles of bunching intensities calculated by pooling all states and years. Intensity is proxied by the percentage of all EITC-eligible taxpayers that report earned income within $500 of the refund-maximizing kink in a given state. Dollar values are adjusted to 2014 levels. The District of Columbia is excluded. This figure was created by the authors using data from the IRS Compliance Data Warehouse.

Figure A9: Bunching intensity at the refund-maximizing kink: 2014

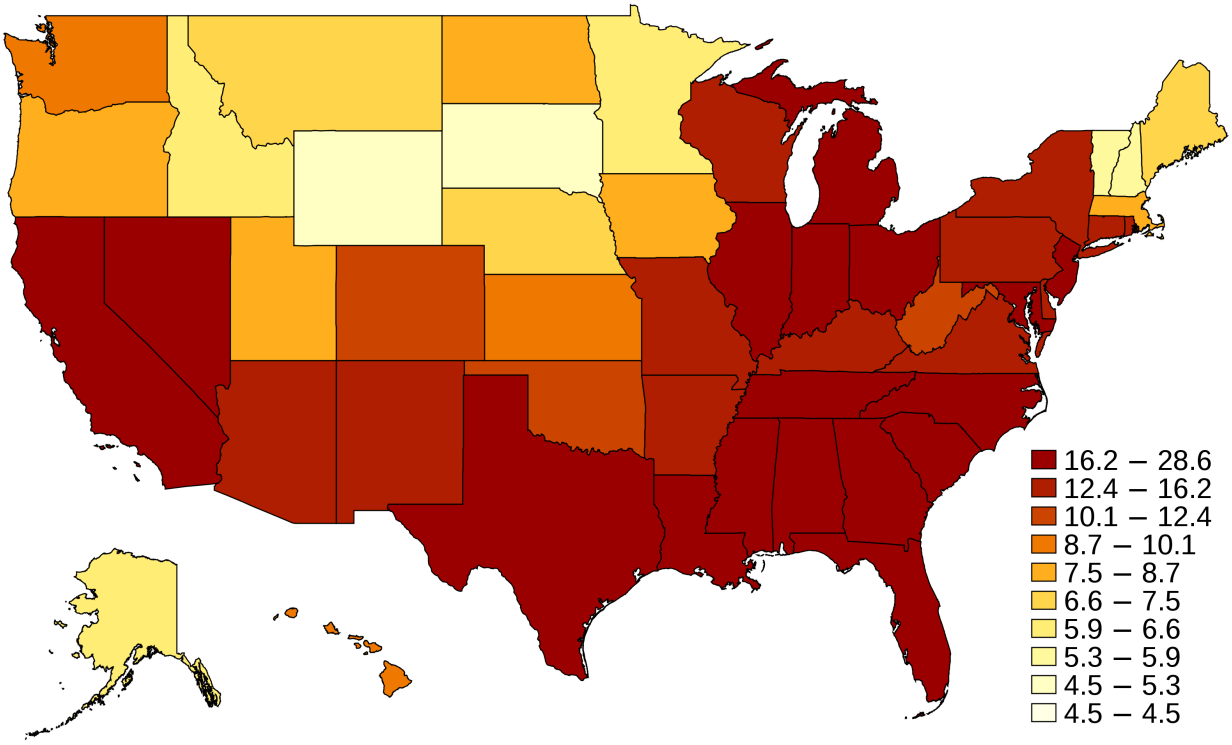

The sample is limited to taxpayers with children and self-employment income. Colors represent deciles of bunching intensities calculated by pooling all states and years. Intensity is proxied by the percentage of all EITC-eligible taxpayers that report earned income within $500 of the refund-maximizing kink in a given state. Dollar values are adjusted to 2014 levels. The District of Columbia is excluded. This figure was created by the authors using data from the IRS Compliance Data Warehouse.