Taxpayer Confusion: Evidence from the Child Tax Credit*

Naomi Feldman† Laura Kawano‡ Peter Katuščák§

August 2014

Abstract

We develop an empirical test for whether households understand or misperceive their tax liability changes. Our identifying variation comes from the loss of the Child Tax Credit when a child turns 17. Using this age discontinuity, we find that despite this tax liability increase being lump-sum and predictable, households reduce their reported labor income when they discover they have lost the credit. This finding suggests that households misinterpret at least part of this tax liability change as an increase in their marginal tax rate. This evidence supports that tax complexity can cause significant confusion and leads to unintended behavioral responses.

*We thank Kate Antonovics, Eli Berman, Sebastian Bradley, Mike Christian, Julie Cullen, John Diamond, Libor Dušek, Jacob Goldin, Roger Gordon, Štěpán Jurajda, Sara LaLumia, Jeff Liebman, Byron Lutz, Olivia Mitchell, Marco Manacorda, Karen Pence, William Peterman, Todd Pugatch, Daniel Reck, Dan Shaviro, Joel Slemrod, and numerous seminar and conference participants for helpful comments and suggestions. Financial support from the Grant Agency of the Czech Republic (402/10/2085) and Global Development Network (RRC IX-04) is gratefully acknowledged. The views expressed in this article are those of the authors and not necessarily reflect those of the Federal Reserve Board or the U.S. Department of Treasury. Any remaining errors are our own. This paper was previously circulated under the title "Effects of Predictable Tax Liability Variation on Household Labor Income."

†Research Division, Federal Reserve Board, Washington, D.C. 20551, USA, naomi.e.feldman@frb.gov.

‡Office of Tax Analysis, U.S. Department of Treasury, Washington, D.C. 20551, USA, Laura.Kawano@treasury.gov.

§CERGE-EI, P.O.Box 882, Politickýchvězňů 7, 111 21 Praha 1, Czech Republic, Peter.Katuscak@cerge-ei.cz.

A fundamental assumption in public finance is that individuals consider taxes when making economic choices. Indeed, a voluminous literature shows that taxes significantly influence behaviors along several margins including labor supply, portfolio allocations, and savings.[1] A standard assumption holds that individuals understand the tax schedule they face. Therefore, the interpretation of this literature is that these behavioral responses arise from changes in tax rates.

However, the U.S. tax code is highly complex. As a result, it may be difficult for taxpayers to understand fully the tax rates that they face and, thus, their tax-induced incentives.[2] The costs associated with learning the details of tax provisions, including cognitive effort, time, and money, may lead some taxpayers to be unaware of their current tax situations or of how their tax liabilities might change in the future. Because the costs of fully understanding the tax system are high, taxpayers may instead use rules of thumb (Liebman and Zeckhauser 2004) or use their interactions with the tax system to learn about the underlying parameters that affect them the most.

In this paper, we devise a test for whether households understand or misperceive changes in their tax schedule. We identify a provision in the tax code that generates a predictable, lump-sum change in tax liability. This variation comes from an age-discontinuity in eligibility for the Child Tax Credit (CTC). To qualify for the CTC in a given year, a child must be younger than 17 at the end of that year. For example, a household with a child who turns 17 in December of year t is not eligible for the CTC for that child for that year; a household with a child who turns 17 in January of the following year is eligible for the CTC for that child for year t (but not for year $t+1$). This credit loss generates variation in tax liabilities that is lump-sum, predictable in advance, and plausibly exogenous. We examine how this

[1] See, for example, Eissa (1995) or Eissa and Liebman (1996) for labor force participation of women, Looney and Singhal (2006) for the intertemporal elasticity of labor earnings, Goolsbee (2000) for the timing of income realization, Poterba and Samwick (2003) for risk-taking and portfolio behavior, and Feldstein (1995), Auten and Carroll (1999), Gruber and Saez (2002), and Kopczuk (2005) for reported and taxable income. On the other hand, Saez (2004) finds that only the top 1% of incomes show evidence of behavioral responses to taxation.

[2] See, for example, Abeler and Jäger (2013) and Chetty, Friedman and Saez (2013).

variation affects reported household labor income.

To motivate our empirical strategy, we develop a model of labor supply choice in the presence of potential misperceptions of the tax system. We identify three types of taxpayers that differ in their beliefs about their current and future tax schedules. Our model generates predictions over the effect of the loss of the CTC on labor supply that differs by type. First, *fully informed* parents anticipate the credit loss in advance and so, barring liquidity constraints, should not adjust their labor supply. Second, *ex post informed* parents fail to anticipate the CTC loss, but understand that they experienced a lump-sum credit loss, *ex post*. These households may increase their labor supply in response due to a non-negative income effect. Lastly, *ex post confused* parents fail to anticipate the credit loss and additionally misinterpret why their tax liability increased. In particular, this type attributes at least some of the tax liability increase to an increase in their marginal tax rate (MTR). If the substitution effect dominates, *ex post* confused households will decrease their labor income in response to losing the CTC. Under our model, a negative labor income response can only be supported by households misperceiving at least some of the credit loss as an increase in their MTR on average.

We construct a panel from population-level U.S. federal tax returns filed between 2004 and 2011. We identify all married-couple households with a child who turns 17 during January or December during this period. Using a regression-discontinuity (RD) design, we find that upon losing the CTC, households reduce their reported labor income by approximately 0.8% relative to households who have just retained the credit for another year. This result is obtained even though losing the CTC has no mechanical impact on MTRs for the selected sample. We show that this effect is not driven by a direct effect of child aging or a spurious correlation between the timing of birth and income growth. We further show that the CTC has no significant effect on reported labor income in the years prior to its loss, suggesting that the effect is not driven by a strategic re-timing of income. We thus interpret our result as evidence that taxpayers confuse part of the credit loss as due to an increase in their MTR.

To understand better the magnitude of our estimated treatment effect, we interpret it within an elasticity framework. This exercise is complicated because it requires assumptions about households' beliefs about the MTR that they actually face prior to the credit loss and the amount of credit loss misperceived as due to an change in MTRs. As such, we do this elasticity calculation under a number of alternative assumptions. First, we assume that households interpret changes to their average tax rate (ATR) as changes in their MTR (i.e., the case of "ironing" in Liebman and Zeckhauser (2004)). This implies a wage income elasticity of 0.4 and an elasticity of taxable income of 0.8. We then generalize this intuition and allow households to perceive their tax liability increase as due to a mixture of an increase in their MTR and a lump-sum change in liability. Using the generally accepted range of elasticity estimates in the literature, we find support for households confusing at least half of the CTC loss as due to an increase in their MTR.

Our paper contributes to a broader literature on taxpayers' misperceptions of the tax system, which we view as encompassing the oft-intertwined issues of tax complexity and salience. Elements within the tax code often interact to create complex structures that may be difficult to decipher.[3] For example, evidence from surveys reveals that taxpayers often are unable to report their correct MTRs (Brown (1968), Fujii and Hawley (1988), de Bartolome (1995) and Romich and Weisner (2000)). Several papers make use of the fact that some taxes are more readily observable (i.e., more salient) than others. The growing literature on tax salience provides ample evidence that individuals are more responsive to sales tax changes that are more visible even if the final tax-inclusive prices are the same (Chetty, Looney and Kroft (2009); Homonoff and Goldin (2013), Feldman and Ruffle (2014) and Feldman, Goldin and Homonoff (2014)).[4] In our setting, the complexity of tax policy changes can lead to confusion over the slope and level of the tax schedule. Just as our results suggest

[3]The income tax system is not the only case where this "price obfuscation" occurs. See, for example, Gabaix and Laibson (2006), Elisson (2006) and Heidhues, K″ oszegi and Murooka (2014) for theoretical models on firms' usage of such techniques.

[4]Similarly, Finkelstein (2009) studies salience in the context of electronic toll debiting. Hossain and Morgan (2006) examine salience in shipping and handling for online purchases, and Choi, Laibson and Madrian (2008) examine salience in fees on index funds.

that a lump-sum change in tax liability is misinterpreted as partly due to an increase in the slope of tax schedule, we may similarly anticipate that policies that change MTRs are partially misinterpreted as a level shift in tax schedule. Generally, tax policy changes that are misunderstood may lead to unintended behavioral consequences, with implications for welfare.[5]

Our empirical framework could be adapted for other settings. For example, one could use other child age-relevant discontinuities, such as the Earned Income Tax Credit (EITC) or child deduction that occur at age 19 (or 24 if a full-time student). These experiments, however, are less than ideal. In both cases, a child aging out of eligibility could trigger an actual change in MTRs. While such changes may be subject to much confusion, separating responses to actual changes in the MTR from perceived changes would be empirically challenging.[6] The age-discontinuity for CTC eligibility provides a clean, natural experiment for testing taxpayer perceptions, and the lessons from this setting provide insights into the effect of confusion over myriad tax provisions, including the loss of the EITC, health care subsidies, and other provisions of the tax system.

1 Model

We formalize a model of labor supply choice in the presence of potential tax schedule misperceptions. Suppose that an infinitely-lived household faces a linear tax schedule in every period s with MTR given by τ_s and demogrant (a lump-sum deduction in tax liability) given by D_s. A household with taxable income y_s thus faces a tax liability of $T_s(y_s) = -D_s + \tau_s y_s$.[7] We assume that the parameters τ_s and D_s are knowable in principle, but allow for households

[5] Several recent papers examine the welfare implications of tax salience, including Chetty et al. (2009), Goldin (2013), and Reck (2013).

[6] A focus on the EITC would likely also include single head of households whose labor supply decisions more complicated. In addition, the impact on MTRs would depend on the number of children. One may argue that an EITC household on the flat portion of the EITC may face a lump sum change upon a child aging out of eligibility but this range is small and would also rely upon lining up the plateau regions of the EITC schedules for one, two or three children.

[7] In practice, income tax schedules are piecewise linear, so the proposed schedule can be thought of as a local approximation of a more complicated tax schedule in the relevant range.

to be unwilling to invest the resources necessary to learn them.[8] Households file their tax return for year s at the beginning of year $s+1$. At this point, households are aware of $T_s(y_s)$, though not necessarily the underlying tax parameters that determined it.

We define three types of households: *perfectly informed*, *ex post informed*, and *ex post confused*. Perfectly informed households know perfectly their current tax schedule. Moreover, these households correctly perceive their past tax schedules and predict future tax schedule changes. In contrast, *ex post* informed and *ex post* confused households do not necessarily know their current tax schedule. We assume that both types use their beliefs about their most recent tax parameters, τ_{t-1} and D_{t-1}, as their expectation for their current and all future tax schedules.[9] Thus, both types fail to anticipate future tax schedule changes.

We distinguish the latter two types by their *ex post* learning when their tax schedule changes. *Ex post* informed households perfectly learn about their period $t-1$ tax schedule upon filing their tax return in period t. This learning could be rationalized by households paying attention to their tax parameters when calculating their tax liability. As a result, in period t, *ex post* informed households have perfect knowledge of all previous tax schedules. Given our previous assumption on expectation formation, these households use this new tax schedule as its expectation over future tax schedules.

In contrast, when filing their period $t-1$ tax return, *ex post* confused households note only that their tax liability differs from their expectations, rather than learning their new tax schedule. We assume that these households use a simple updating rule in which some fraction α of the surprise is attributed to an MTR innovation and the remaining fraction $1-\alpha$ is attributed to a demogrant innovation.[10] Denoting τ_s^t, D_s^t, and T_s^t to be period t's

[8] At the cost of a more complicated exposition, surprises to the current period's tax schedule could be incorporated in the analysis as well.

[9] This simple adaptive process of expectation formation satisfies the Law of Iterated Expectations and hence is internally consistent in a minimal way. For the purpose of the empirical prediction, this assumption can be relaxed somewhat as long as a higher realization of $T_{t-1}(y_{t-1})$ increases the expectation of future MTRs and reduces the expectation of future demogrants among the *ex post* confused households, each by the same amount in the current and all future time periods.

[10] If we assume that the household is only confused about the MTR, but not about the intercept, an unexpected tax liability increase will increase the expected future MTRs by the magnitude of the surprise in the realized ATR. If coupled with the assumption that there is a demogrant, this case corresponds to the

beliefs of τ_s, D_s, and T_s, for all $s \geq t$,

$$\tau_s^t = \tau_{t-1}^{t-1} + \alpha \frac{T_{t-1}(y_{t-1}) - T_{t-1}^{t-1}(y_{t-1})}{y_{t-1}}$$
$$D_s^t = D_{t-1}^{t-1} - (1-\alpha)\left\{T_{t-1}(y_{t-1}) - T_{t-1}^{t-1}(y_{t-1})\right\}. \tag{1}$$

Importantly, these households typically misunderstand the tax schedule changes they experience.

Households maximize their discounted utility defined over consumption and leisure. In period t, households solve the following intertemporal problem:

$$\max_{\{c_s, l_s\}_{s=t}^{\infty}} E_t \left[\sum_{s=t}^{\infty} \delta^{s-t} u_s(c_s, 1-l_s) \right] \tag{2}$$

$$\text{s.t. } a_t \text{ given} \tag{3}$$

$$\text{s.t. } a_{s+1} = (1+r_s)\left[a_s + (1-\tau_s)w_s l_s + D_s - c_s\right] \text{ for all } s \geq t \tag{4}$$

$$\text{s.t. } \lim_{s\to\infty} a_s \prod_{u=t}^{s-1} (1+r_u)^{-1} \geq 0. \tag{5}$$

In this formulation, $u_s(\cdot)$ is the period-s flow utility function defined over consumption (c) and leisure ($1-l$), l is labor supply, $\delta \in (0,1)$ is the per-period discount factor, a_s is the asset stock at the beginning of period s, r_s is the net rate of return on assets between periods s and $s+1$ and w_s is the wage rate in period s.[11] The operator $E_t(\cdot)$ captures type-specific expectations about current and future tax schedules. We assume that both consumption and leisure are weakly normal goods.[12]

"ironing" hypothesis in Liebman and Zeckhauser (2004). Under ironing, a household predicts its current MTR is the previous period's ATR.

[11] For simplicity, we assume that the rate of return and the wage rate processes are non-stochastic. This assumption is not essential for the model and could be relaxed at the cost of increased expositional complexity. The constraint (5) is the usual transversality condition. If the household is subject to a credit constraint, (5) is replaced by $a_s \geq 0$ for all $s \geq t$.

[12] With tax rate misperceptions, households will exhaust their perceived budgets given their beliefs and, as a result, there will generally be a difference between the expected and actual tax liability. To resolve this discrepancy between actual expenditures and the available budget, previous research has defined a budget adjustment rule as a specification of which good bears some adjustment necessary to equate expenditure to the available budget (Reck 2014). We implicitly assume a zero budget adjustment effect on current

To motivate our empirical strategy, consider the effect of a permanent decrease in the demogrant in period t on household labor supply choices.[13] Fully informed households anticipate this change and had incorporated it into their lifetime optimization problem. In the absence of credit constraints, there is no impact on household labor supply.[14] *Ex post* informed households are surprised by a higher than expected period t tax liability upon filing their tax return in period $t+1$, but they perfectly understand that the increase was due to a decrease in the demogrant. They also expect that the demogrant has decreased from that period onwards. As a result, there is a negative wealth effect that (weakly) increases labor supply. In contrast, *ex post* confused households are surprised by a higher than expected period t tax liability but, in line with (1), form a higher expectation of $\{\tau\}_{t-1}^{\infty}$ and a lower expectation of $\{D\}_{t-1}^{\infty}$. The perceived increase in their MTR generates both a negative substitution effect that decreases labor supply and a negative wealth effect, while the decrease in the demogrant generates a negative wealth effect that increases labor supply. The overall impact on labor supply depends on the relative magnitudes of these effects. Put together, our model predicts that these agent types respond differently to an unexpected decrease in the demogrant. In particular, households who reduce their labor supply in response must be confused and have a dominant substitution effect.

2 Identification Strategy and Empirical Implementation

To test how taxpayers interpret changes in their tax liability, we identify a source of predictable, lump-sum variation in after-tax income. This variation comes from an age-based

consumption and labor supply in period t. That is, any discrepancy between the realized and expected tax liability in period t is projected to a_{t+1} so future consumption and labor supply bear the full burden of budget adjustments resulting from tax schedule misperceptions. We believe this is a natural assumption to make in the context of dynamic optimization.

[13] Similar thought experiments for changes to the MTR can easily be conducted.

[14] For all types, a binding credit constraint results in households consuming less and working more because of an inability to smooth consumption and leisure using future income.

discontinuity generated by the eligibility rules for the Child Tax Credit (CTC). The CTC was introduced in 1998 and provides a non-refundable credit for an eligible child below 17 years of age as of December 31 of the tax year.[15] The CTC amount was initially set at $400 per eligible child and has increased over time. We focus on the recent period (2004-2011) when the CTC was $1,000 per eligible child.[16] In addition to the CTC, the Additional Child Tax Credit (ACTC) was introduced, which provided limited refundability of the non-refundable part of the CTC for families with three or more qualifying children.[17] In 2001, the ACTC was expanded to allow any family to claim the non-refundable part of the CTC up to one tenth of the excess of their earned income over $10,000.[18] The CTC is phased out with adjusted gross income (AGI) at a 5 percent rate above $110,000 for married couples filing a joint tax return[19]

Four features of the CTC make it a good natural experiment for our analysis. First, to be eligible for the credit, the dependent child must not have reached 17 years of age by December 31 of the tax year in question. Because a child's 17th birthday is perfectly predictable, so should be the associated net income loss. Second, over the period we consider, virtually any household with AGI between $30,000 and $100,000 can take advantage of the full $1,000 of the CTC because of the ACTC. As a result, the loss of the CTC constitutes a pure lump-sum change in both tax liability and after-tax income conditional upon claiming

[15] There are several provisions in the tax code that make the tax schedule a function of a dependent child's age, such as the loss in the eligibility for the personal exemption and the EITC for a dependent child who turns 19 (or 24, if a full time student). This provision has been exploited by Looney and Singhal (2006) and Dokko (2007) in order to estimate the effect of marginal tax rates on labor supply.

[16] The CTC was increased to $500 for the 1999 and 2000 tax years, to $600 for the 2001 and 2002 tax years, and then to $1,000 for the 2004 tax year, where it has remained since then. 2003 was a slightly unusual year where the credit transitioned from $600 to $1,000. Upon filing taxes, the credit was $600, but then an additional $400 lump sum was provided retroactively. In addition, as part of stimulus payments in 2008, eligible households received an additional $300 on top of the $1,000 credit.

[17] These families could claim the non-refundable part of the CTC up to the amount of employee contributed Social Security and Medicare taxes less any EITC they received.

[18] The $10,000 threshold has been indexed to inflation over time. In addition, starting in 2004, the ACTC limit was increased to 15 percent of earned income in excess of the threshold. Families with three or more eligible children could still claim the non-refundable part of the CTC up to the amount of employee contributed Social Security and Medicare taxes less any EITC they received if this limit turned out to be higher.

[19] The thresholds are $75,000 and $55,000 for single/head of household taxpayers and married taxpayers filing separately, respectively. None of these thresholds are indexed for inflation.

the dependent. Third, there is no other age-dependent credit that is lost upon turning 17. Finally, it is difficult to plan the timing of birth. Thus, among families whose children turn 17 just before the end of year t or at the very beginning of year $t+1$, eligibility for the CTC is virtually exogenous.[20] Focusing on households whose children turn 17 around the turn of a year, losing the CTC generates exogenous, predictable, and lump-sum variation in tax liability.

The loss of the $1,000 CTC is at least as large as other empirically analyzed tax-related income shocks. For example, the income tax rebates of 2001 were $600 for married couples filing joint returns, and the average tax refunds analyzed by Souleles (1999) were $874 based on Consumer Expenditure Survey data. Accounting for inflation would make the loss of the CTC and these refunds fairly comparable. When households in our sample are still CTC-eligible, the average refund amount is nearly $2,800. This would make the CTC worth over one-third of the average refund. The comparison of the CTC to the EITC is less straightforward as the EITC applies to a different population than the one studied here and is more highly dependent upon marital status and number of children. Nonetheless, for a married couple earning the maximum amount under the EITC, the impact of a child aging out of EITC-eligibility (say, going from two children to one) is, in our most generous calculation, slightly less than twice as large as the impact of losing the CTC during the 2004-2011 time frame.

We designate treated households as those whose child turns 17 in December of year t, thus making that child ineligible for the CTC in that year. If unaware that their child will lose CTC-eligibility, these treated households are unlikely to notice the CTC loss until the early months of year $t+1$ when they file their tax returns, even though they have technically lost the credit for tax year t. The control group consists of households whose child turns 17 in January of the following year. Thus, these children just retain CTC-eligibility for tax year t and instead become ineligible in year $t+1$. These control households are similarly

[20] In section 4, we discuss evidence why this may not be the case and test the robustness of our results to the potential endogeneity in the timing of birth.

unlikely to notice the credit loss until year $t+2$ if they are unaware that their child becomes ineligible for the CTC. Our identification strategy thus relies on households who lose the credit at the end of one tax year being similar to households who just retain the credit for that same year but for the loss of the CTC.

To test whether households are *ex post* confused by the loss of the CTC, we examine the effect of realizing this loss on household labor supply choices. The natural experiment that we exploit forms the basis for employing a difference-in-differences (DID) methodology. Further, the age-eligibility rule for the CTC allows us to use a regression discontinuity (RD) methodology within the DID framework. We define our running variable, N_days, to be the number of days until January 1 of year $t+1$, so a household is assigned to the treatment group when $N\_days < 0$.

The baseline equation that we estimate is given by:

$$\Delta \ln Y_{it+1} = \beta_0 + \beta_1 T_{it} + \pi' X_{it} + f(N\_days) + \gamma_t + u_{it+1}, \tag{6}$$

where Y is household wage income, X is a vector of household characteristics in year t, and γ_t are year fixed effects. The variable T is an indicator for whether the household's child has turned 17 in December of year t, and $f(N\_days)$ is a polynomial expansion of the number of days. We allow the parameters of $f(\cdot)$ to vary on either side of the discontinuity. The DID estimator controls for unobserved heterogeneity that is time invariant, which is particularly important in our setting due to the very limited demographic information available in administrative tax data.[21] The vector X_{it} contains the age of the primary filer (the level and its square), the number of children in the household (the level and its square), and state of residence, each in year t. In principle, if the treated and control households are identical but for the loss of the CTC, a comparable model with the dependent variable in year $t+1$ levels should provide similar results. We return to this specification in Section 4.3.

[21] An alternative approach would be to estimate a fixed effects specification with $\ln Y$ as the dependent variable. With two time periods, as in our context, these two approaches are equivalent.

The null hypothesis is that households *ex post* understand the tax schedule that they face. Perfectly informed households anticipate the tax schedule change *ex ante*, so there should be no reaction in labor supply. *Ex post* informed households are surprised by the increase in tax liability but correctly perceive the tax liability change as due to the loss of the credit *ex post*. In this case, there is a non-negative income effect to labor supply. In either case, an *ex post* understanding of the tax liability change implies that, combined with a possible tightening of liquidity constraints, there is a non-negative impact on parental labor income in the subsequent year ($\beta_1 \geq 0$). The alternative hypothesis is that households are *ex post* confused. In this case, households both fail to anticipate the loss of the credit and misinterpret the change in tax liability as at least partially reflecting an increase in their MTR. If the income effect dominates, then $\beta_1 \geq 0$ and we are unable to distinguish between the null and alternative hypotheses. However, if $\beta_1 < 0$, it must be the case that, on average, households are *ex post* confused and that the substitution effect dominates the income effect.

3 Data

Our data come from administrative, population-level U.S. federal tax filings.[22] We first identify all individuals who turn 17 between 2004 and 2011 in either December or January.[23] We match these children to the primary filer who claims the child as a dependent on his tax return, requiring that the same primary filer claims a child across years to avoid changes in income resulting from changes to household composition.[24]

[22]These data are housed at the IRS's Compliance Data Warehouse (CDW), which contains transcribed data from all individual tax returns and information returns that are filed since 1999, along with Social Security administrative data.

[23]These individuals are identified using date of birth information from Social Security administrative records. There are roughly 400,000 children born in each month over our sample period. We begin our analysis in 2004 to consider the period when the CTC is $1,000 per eligible child, rather than have household responses differ over time due to different credit amounts.

[24]Family relationships cannot be determined from tax return data, and so we consider these matched primary filers to be a parent. We drop the few instances where the dependent is claimed as a spouse. Note that children in treated households turn 17 and 18 in years t and $t+1$, whereas children in control households turn 16 and 17 in these years. We restrict our analysis to households who consistently claim these dependents in the years that they turn 16, 17 and 18 years old. There are several reasons why kids may not be claimed

We collect income variables for the relevant tax years, along with the limited demographic information available from tax returns (i.e., the number of children, age of the primary filer, and filing status). We make several additional sample restrictions. First, we restrict our analysis to married couples filing joint returns so we do not conflate changes to labor income due to the change in the CTC with those due to changes in marital status.[25] Second, we exclude households where the primary is younger than 35 years old (which would put the primary filer at around 18 years of age at the time of birth of the child) or older than 60 to remove the potential influence of completing education or retirement decisions on labor supply choices. Third, we exclude households outside of the 50 states or D.C. We also exclude households in states where it is likely that treatment and control children are in different academic grades.[26] We also exclude roughly 1% of households with more than one child turning 17 in a particular month-year. Finally, we focus on households who plausibly face a pure lump-sum tax liability change due to their child aging out of CTC eligibility. This restriction comes in two forms. First, we restrict attention to households whose incomes place them beyond the refundable portion and before the phase-out region of the CTC. That is, we retain households with AGI between \$30,000 and \$100,000, using reported AGI in year $t-1$ to avoid endogenous income choices in the treatment year. Second, we compare the actual MTR that treated households faced in year t with the hypothetical MTR they would have faced had their child not aged out of CTC-eligibility. We exclude treated households whose loss of a CTC-eligible child generated a change in their MTR.[27] All dollar values are winsorized at the 1% and 99% levels to mitigate the influence of outliers.

Our final estimation sample includes 851,890 observations. We construct seven pair-wise

consistently over time. Some parents will become non-filers and there may be changes in who claims a child. Some fraction of kids die before age 18 and some may become primary filers.

[25] In general, IRS statistics suggest that over 95% of married households file jointly.

[26] We exclude households from five states that had a December 31 school entry cutoff date in the year in which the child was born: Delaware, Hawaii, Louisiana, Maryland and Rhode Island. In these states, there is likely a one academic year difference between the treatment and control groups. Note, however, that we have the state in which the household filed, which is not necessarily the state in which the child lived when he or she first entered school.

[27] We compute these MTRs using the NBER's TAXSIM calculator (Feenberg and Coutts 1993). This final restriction affects 9% of the remaining sample of households.

cohorts of treated households (i.e., those whose kids turn 17 in December of year t) and control households (i.e., those whose kids turn 17 in January of $t+1$). Figure 1, which plots the density of households by day of birth of the child under consideration, shows that the distribution of households is fairly smooth, with expected dips in births around Christmas and New Year's.

Figure 2 provides an initial verification of our experiment and some preliminary evidence of taxpayer confusion. The top panel of Figure 2 shows the sharp discontinuity in CTC-eligibility. This figure plots the average CTC amount claimed for our treatment and control groups averaged over day of birth around the discontinuity.[28] We include a cubic polynomial of the running variable on each side of the discontinuity.[29] As expected, CTC amounts are roughly the same in year $t-1$, when children in both the treated and control groups are eligible for the CTC. In year t, children in our treated group have just become ineligible for the CTC, whereas the control group's children has just remained eligible for an additional year of credit. There is a sharp discontinuity in the amount of CTC claimed, with the differential in credits being roughly $1,000. In year $t+1$, children in the control group have now also become ineligible for the CTC, and there is again no difference in the CTC amounts claimed by the treated group and control group.[30]

The bottom panel of Figure 2 provides some evidence of the treatment effect we would expect if households misunderstand the tax liability increase experienced due to the loss of the CTC. This figure depicts log wages for our treatment and control groups. Household wage income looks very similar across treated and control households in year $t-1$, the last year in which both sets of households receive the CTC. In year t, treated households have

[28] A related concern is that our sampling methodology identifies intent to treat rather than treatment on the treated. In practice, there is very little difference between these two estimators, as take-up of the CTC is extremely high. For example, over 97% of households in our control group report some CTC in year t, the last year in which they are eligible to claim the credit for the child that we consider.

[29] These figures look very similar using different specifications of the running variable.

[30] Even if ineligible for the CTC, children can still be claimed as dependents for purposes of the EITC, dependent exemptions, and head of household status. In figures not shown (available upon request), we verify that the treatment effect is not due to a differential change in the number of children claimed as dependents between the treatment group and control group. Thus, changes in CTC amounts are not driven by other children being disproportionately removed from treated households.

lost the CTC for that tax year but do not appear to exhibit any change in their wage income relative to control households. Household wage income does not show evidence of a change at the discontinuity. In the early months of year $t+1$, however, treated households would see their change in tax liability due to the CTC when they filed their tax returns. We do see evidence of a discontinuous break in household wage income in this year, suggesting that this tax liability change was both unanticipated and interpreted as related to a change in marginal tax rates.

Table 1 provides summary statistics for our sample by treatment status. Overall, treated and control households appear very similar in both years t and $t+1$ for variables that are not affected by the loss of the CTC. Treated and control households are balanced on observable demographic characteristics (i.e., number of children, age of the primary filer), however due to our large sample sizes, each is statistically different at the 1% level. Importantly, for those variables where there is a meaningful difference between treated and control households in year t, they appear to be driven primarily by the loss of the CTC. Treated households have \$923 less in CTC amounts on average in the year that their child becomes ineligible for the credit. This treatment is reflected in average tax liabilities: tax liabilities for the treated group are roughly \$760 higher than those in the control group in year t, but this difference practically disappears the following year.[31] The difference in tax liabilities does not appear to be driven by a difference in EITC amounts. There is also evidence that treated households likely would have noticed the loss of the CTC in the early months of year $t+1$ when they filed their taxes, as seen in the change in refund propensities and refund amounts. On average, in year t, treated households are 8% less likely to receive a refund and have \$826 less in refunds than those in the control group.

The last three rows of Table 1 provide information on certain income sources. The first is our primary dependent variable of interest: wage and salary income. This income amount is reported on line 7 on the Form 1040 and is highly accurate because it is verified by W-2s filed

[31] Tax liabilities are computed as the total tax owed, less EITC and ACTC payments.

by employers with the IRS.[32] The next variable is a broader measure of labor income that includes self-employment earnings reported on Schedule C. Self-employed individuals exhibit higher levels of tax evasion and avoidance (Slemrod 2007), so changes in this broader labor income measure may be influenced by the income reporting decisions of the household.[33] Lastly, we report the household's taxable income, which includes all forms of income minus deductions. Changes to taxable income will capture a wide array of behavioral responses, including changes to labor supply, savings, and tax evasion or avoidance behaviors.

4 Results

4.1 Main Result

Results from our estimation of equation (6) are presented in Table 2, where the dependent variable is the log change in household wage and salary income. Each column incrementally increases the order of the polynomial of the running variable (days from January 1). Across all specifications, the point estimates suggest that upon losing the lump-sum credit, households reduce their wage income. This negative effect on labor income is consistent with at least some households suffering *ex post* confusion, misinterpreting this lump-sum change as at least partially due to an increase in their MTR. The results are fairly consistent across the board, though sensitive to the order of the polynomial of the running variable. The treatment effect ranges from -0.003 to -0.008.[34] If we take column (4) as our preferred specification, treated households reduce their wage earnings by 0.8% in the year that they observe a $1,000 increase in their tax liability, relative to otherwise similar households who retained the CTC for an additional year.[35] We provide some interpretations of the magnitude of this

[32] Slemrod (2007) shows that approximately 99% of wage earnings that should be reported to the IRS are reported.

[33] Also note that self-employment earnings may be negative.

[34] The estimated treatment effect in column (3) has a p-value of 0.21.

[35] We view the third-order polynomial specification as a preferred balancing of statistical power and overfitting the data.

estimated effect in the next subsection.

In the top row of Table 3, we present results using two alternative dependent variables that differ in their tax reporting requirements and in their available margins for response. Evaluating changes to these alternative forms of income may provide insight into the extent to which households adjust their behavior in response to the perceived change in MTRs. These adjustments may come in the form of changes in labor supply (or effort), the timing of income, or tax evasion and avoidance decisions. Column (1) repeats the finding from column (4) of Table 2 for ease of comparison.

In column (2), we expand the dependent variable to a broader measure of labor income that includes self-employment income. Self-employment earnings are easier to manipulate than labor income that is subject to third-party reporting, so we might expect a larger effect of a perceived increase in MTR on reported labor earnings. However, we find that adding self employment income to wages and salaries generates little change in the estimated treatment effect. In fact, in unreported specifications, we incrementally restrict the households that we include in the regression based upon the fraction of labor income that is earned from self-employment. Focusing on households who earn a substantial fraction on their income from sole proprietorships has virtually no impact on the treatment effect. These results suggest that our estimated treatment effect reflects real changes in labor supply and that reporting responses are negligible.

Our dependent variable in column (3) is based off of taxable income. In additional to labor supply responses, this income measure can be influenced by a variety of responses, such as itemized deductions, capital gains realizations, and other credits. Here, we find that the treatment effect is roughly twice as large as that of only wage and salary income suggesting that, in addition to changes in labor supply, households make use of other margins of response. Overall, these results align with more traditional analyses of behavioral responses to real changes in MTRs. Our results further show, however, that households also respond to perceived changes in MTRs even when they have not actually experienced a change in

their MTR.[36]

4.2 Implied Elasticities

The estimated treatment effects indicate that upon losing the CTC some households interpret their tax liability increase as at least partially due to an increase in their MTR. A natural next question is whether the magnitudes of the estimated effects are plausible. However, interpreting these magnitudes requires additional assumptions over the perceived change in MTR to which households react. In this Subsection, we consider a number of scenarios that might describe a household's misinterpretation of the CTC loss and compute the implied elasticity based on the perceived change in MTR and observed change in income. Overall, we find that under a broad range of assumptions, the implied elasticities fall within the range of elasticities in the literature that are estimated based off of real changes in MTRs.

First, we consider the special case when households disregard the demogrant when thinking about the tax schedule. Under this scenario, households "iron" out the tax schedule and confuse their ATR with their MTR.[37] Our estimated treatment effect would then reflect responses to the change in ATRs due to the loss of the CTC. To compute the implied elasticity under "ironing," we calculate the change in ATR that treated households face when they lose the CTC.[38] The uncompensated elasticities based on the implied change in ATR and the estimated treatment effect are presented in the second row of Table 3 The implied elasticities of wage earnings, labor income (wage earnings plus self-employment income) and taxable income are 0.4, 0.5 and 0.8, respectively. It is fairly intuitive that these latter two elasticities are larger than that of wage income as they include elements of income that are less subject to information reporting and are more prone to manipulation (timing or

[36]Relatedly, Miller and Mumford (2014) find evidence of a child care expenditure response to a perceived change in the value of the 2003 expansion of the CTC and little evidence of any response to the actual change.

[37]This particular type of confusion is referred to as "ironing" in Liebman and Zeckhauser (2004)

[38]We first compute the ATR for treated households as the ratio of tax liability and AGI in year $t-1$, the year before they lose the credit. We then compute the ATR implied after then \$1,000 loss based on information in year $t-1$. This is the ATR for households prior to any other behavioral response.

avoidance, for example).

If we consider our wage income elasticity as a proxy for the household labor supply elasticity, our estimates from Table 3 are comparable to previous studies. For example, the elasticity for prime-age single workers ranges roughly between 0.1 and 0.3 (see, for example, Altonji (1986), McClelland and Mok (2012), and Peterman (2013)) whereas the elasticity for married women is around 0.8 (Eissa 1995). Similarly, estimates of the elasticity of taxable income generally range from 0.2 to 0.8 (Saez, Slemrod and Giertz 2012).[39]

Next, we interpret our estimated treatment effect in our more general model of household confusion. As discussed in Section 1, when households experience a change in tax liability they can attribute some fraction α as being due to a change in the MTR and $1-\alpha$ on being due to a change in the demogrant. Without further assumptions we are unable to pin down α. Moreover, calculating the log change in net-of-MTR necessitates taking a stand on base MTRs that households believe they are facing. This is particularly challenging within the context of taxpayer confusion.

To remain agnostic on the degree of baseline knowledge of the tax system and the degree of *ex post* confusion, we present a range of possibilities in Figure 3. The Figure shows five lines that differ by the baseline MTR that the household may believe it faces. The x-axis measures the amount of the \$1,000 tax liability increase that is attributed to a change in MTR. The far right value of \$1,000 represents a household that believes that all of the increase is due to a change in the MTR (as assumed under "ironing") and \$0 represents a household that attributes the entire tax liability increase to a change in the demogrant (a perfectly informed or *ex post* informed household). The lines thus depict the implied elasticity based upon the perceived change in the MTR under various scenarios and our estimated treatment effect. For example, if a household believes that it faced a 30% MTR in year t and attributes \$600 of the increase in tax liability it realizes in year $t+1$ as due to

[39] In one sense, it is difficult to compare our elasticity calculation to those in the previous literature because we work from the assumption that the changes in marginal tax rates that were used to generate the elasticity estimates were perfectly understood. If not, then those elasticities may be miscalculated as well.

an increase in its MTR, the implied wage income elasticity is roughly 0.60.

There are several major points to take away from this graph. First, if we assume that all households attribute the entire $1,000 tax liability increase to an increase in the MTR (i.e. "ironers"), then the perceived baseline MTR does not much matter. Whether households assume 5% or 40%, the difference in the estimated elasticity between these two extremes is approximately 0.2. In other words, at some point, household confusion over complex tax liability changes dominates knowledge over its true MTR. Second, as the weight placed on the MTR (α) increases, the perceived baseline MTR is of greater importance. For example, if $\alpha = 0.1$, or, only $100 of the $1,000 tax liability increase is attributed to an increase in the MTR on average, we obtain elasticities that are well above 2.0. This suggests that this relatively low level of confusion is not empirically supported given the generally accepted range of estimated elasticities. Finally, if we take the Hicksian aggregate hours elasticity from Chetty (2012) of 0.58 as the "true" elasticity, our calculations imply that households, on average, attribute at least half of the tax liability increase as due to an increase in their MTR.[40] Thus, abstracting from income effects, if our sample consists of "ironers" and otherwise informed households, at least half of all households are confused as to the source of tax liability changes that result from losing the CTC.

4.3 Robustness Tests

In Table 4, we present estimates of a number of extensions and robustness tests. First, while our preferred specification is a difference-in-differences model, estimating a level specification should provide similar results to Equation (6) provided that our treatment and control groups are otherwise identical in the years prior to the loss of the credit. To verify this, columns (1) and (2) present the results using $log(wages)_{t+1}$. These level regressions provide practically identical results to our baseline estimate, though statistical significance is weaker in column

[40] If we assume a 5% base MTR, an elasticity of .58 implies that $500 out of the $1,000 is attributed to a change in the MTR. Analogously, assuming a base MTR of 40% implies that about $750 is attributed to a change in the MTR.

(2) (p-value of 0.25) when we exclude the control for wage income in year $t-1$ that serves as a proxy for unobserved heterogeneity. This is not wholly unsurprising, as the R-squared drops sharply moving from column (1) to column (2) when we exclude this important control, and the standard error of the estimate notably increases due to the loss of precision.

Second, the assumptions underlying our estimating equation imply that we should find no impact of the treatment effect in year t (turning 16 years old). Columns (3) and (4) provide results from estimating a level specification using $log(wages)_t$ as the dependent variable, and column (5) provides results from estimating Equation (6) for year t (i.e., using the difference between years $t-1$ and t). Consistent with our underlying assumptions, log wage income for the treatment group is not significantly different from that of the control group in year t (and is consistent with the bottom middle panel of Figure 2). More importantly, the point estimates are very close to zero.[41] Columns (1)-(5) provide important support to our baseline results and underlying assumptions in Tables 2 and 3. In particular, we can recover the DID estimate as the difference in point estimates from the level specification in year $t+1$ and in year t.

Next, the validity of our identification strategy rests on the assumption that there is no spurious non-tax correlation between the outcome and the treatment variable. Of particular concern is that there may be a direct effect of child aging on parental labor supply. Households in the treatment group have, on average, slightly older children than households in the control group. If the growth rate of parental labor income depends on the age of their children, then our estimated treatment effect may confound the effect of losing the CTC with a direct effect of the child's age. We test for a direct effect of child aging by comparing parental labor income growth rates for households with children turning 17 in consecutive birth months within a tax year.[42] For example, we compare January ("treated") and February ("control") births within the same year. Households in the "treatment" and "control"

[41] Repeating this test on a smaller collected sample in year $t-1$ (15 year olds) also shows no economic or statistical difference between treatment and control. (Unreported but available upon request.)

[42] Because of the intensity of the data collection, for this exercise, we collect data only for all individuals who turned 17 in 2006 (chosen randomly).

groups have children with similar differences in age as our baseline specification, but there should be no difference in the tax consequences between the two groups because they all turn 17 within the same year. We run this placebo test on six cohorts within the year where each cohort is a set of adjoining months (January-February, March-April, etc.). The results, reported in column (6), confirm our hypothesis and show no statistical or economic difference between "treatment" and "control" groups. Thus, our main results do not appear to be driven by a direct effect of child aging on parental labor income nor spurious correlation due to the end of the tax year.[43]

Finally, in columns (7) - (9), we estimate Equation (6) separately for three different income groups: \$30-50K, \$50-70K, and \$70-100K. The treatment effect is strongest for the lowest income category at -0.018 and is statistically significant at the five percent level. The point estimate falls (in absolute value) in each of the remaining two income categories– to -0.008 (p-value of 0.104) and then to 0.003 (not statistically significant at conventional levels). Because the CTC is a fixed amount per child, the effect of the credit loss as a percentage of income decreases in income and, as expected, has a less noticeable impact. This heterogeneity in responses may also reflect that higher income households have a better understanding of the CTC eligibility rules and their effect on MTRs.

5 Conclusion

The complexity of the income tax system can make it difficult for taxpayers to understand the tax schedule that they face. We present a model in which households may misperceive the nature of tax policy changes that they experience: fully informed households fully an-

[43] Additionally, Dickert-Conlin and Chandra (1999) argue that if a child is due around the turn of the year, parents may have a preference to accelerate the birth on the margin so that they can claim tax benefits for the ending calendar year. The authors also find that such behavior is more prevalent among higher income households, raising another potential spurious correlation problem. However, this problem does not pose a concern for our interpretation of the results since it implies that we tend to underestimate β_1 in absolute value as more sophisticated taxpayers would be less subject to confusion, suggesting that the true coefficient is even more negative. Further, more recent analysis by LaLumia, Sallee and Turner (2014) using the universe of tax returns finds that such an effect is quite small.

ticipate predictable tax policy changes; *ex post* informed households fail to anticipate tax policy changes but correctly perceive them once they are experienced; and *ex post* confused households both fail to anticipate tax policy changes and remain confused as to the nature of the policy *ex post*. This latter case generally leads to *ex post* misperceptions over changes in the MTR the household faces. We examine the misperception hypothesis by measuring taxpayer labor income responses to a source of exogenous, lump-sum and predictable variation in the tax liability due to losing eligibility for the Child Tax Credit when the child turns 17.

We find that, on average, households who lose the credit due to having their child turn 17 at the end of a calendar year report about 0.8% lower household labor income in the subsequent year compared to households that have their child turn 17 at the beginning of the following calendar year. This result is inconsistent with all households being fully or *ex post* informed about predictable changes in their tax liability. We argue that at least some households misinterpret the increase in their tax bill as (at least partly) due to an increase in their MTR, leading to a reduction in labor income due to conventional substitution effects. This finding is robust to a variety of tests that include placebo effects at various other age and calendar cutoffs.

Our results suggest that tax policy changes that are not well-understood by the affected population may have unintended behavioral and welfare consequences. In particular, changes that affect the level but not the slope of the tax schedule may result in unanticipated (by policy makers) substitution effects, hence increasing or reducing the deadweight loss relative to the full-information case. On the other hand, changes that mostly affect the MTR may be partly interpreted as changes in the level of the tax schedule, with analogous implications for deadweight loss. The complexity of the tax system may therefore interact with tax changes to create departures from conventionally understood welfare effects.

Table 1: Summary statistics

	Year of Treatment (Year t)		*Year after Treatment (Year t+1)*	
	December birth (T=1)	January birth (T=0)	December birth (T=1)	January birth (T=0)
Age of primary filer	46.4	46.5	47.4	47.5
	(5.59)	(5.53)	(5.59)	(5.53)
Number of children	2.26	2.25	2.18	2.18
	(1.06)	(1.06)	(1.06)	(1.05)
Child Tax Credit (CTC+ACTC)	873	1796	856	842
	(980)	(1,010)	(983)	(980)
Tax liability	4,812	4,054	5,212	5,137
	(4,179)	(4,249)	(4,688)	(4,810)
EITC amount	129	117	168	159
	(568)	(539)	(663)	(647)
EITC, > 0	1,737	1,699	1,884	1,882
	(1,239)	(1,230)	(1,306)	(1,302)
Proportion with a refund	0.79	0.87	0.79	0.79
	(0.40)	(0.33)	(0.40)	(0.41)
Refund amount	1,970	2,796	2,064	2,014
	(3,238)	(3,205)	(3,226)	(3,295)
Average tax rate	7.08	6.82	7.12	6.89
	(3.24)	(3.43)	(3.44)	(3.59)
Wage and salary income	64,811	65,494	66,664	67,494
	(23,701)	(23,850)	(26,327)	(26,494)
Wage and self-employment income	66,677	65,576	68,356	67,425
	(25,335)	(26,492)	(22,940)	(24,329)
Taxable income	40,926	39,945	43,016	42,172
	(27,956)	(29,745)	(35,517)	(45,499)
Number of observations	422,817	429,073	422,817	429,073

Children with December birthdays turn 17 (i.e., households lose the CTC) and those with January birthdays turn 16 in year t.
Source: U.S. federal tax filings housed at the IRS's Compliance Data Warehouse, 2004-2011.

Table 2: Pooled RD Specifications

	Dependent variable: $\Delta \log(WageIncome)_{t+1}$				
	(1)	(2)	(3)	(4)	(5)
T	-0.006***	-0.004***	-0.003†	-0.008**	-0.008*
	(0.001)	(0.002)	(0.002)	(0.003)	(0.004)
Running variable order	0	1	2	3	4
R-squared	0.007	0.007	0.007	0.007	0.007

Control variables (estimates not displayed) are age and age squared of the tax filer, total number of dependent children in the household and its squared term, state fixed, and year fixed effects. Observations: 851,890. Robust standard errors, clustered by household, are in parentheses. *** p<0.01, ** p<0.05,* p<0.1, † p<0.21.

Table 3: Alternative Dependent Variables and Elasticities (Pooled RD)

Dependent variable:	$\Delta \log(WageIncome)_{t+1}$	$\Delta \log(LaborIncome)_{t+1}$	$\Delta \log(TaxableIncome)_{t+1}$
	(1)	(2)	(3)
T	-0.008**	-0.009***	-0.015***
	(0.003)	(0.003)	(0.005)
Elasticity	0.44	0.47	0.81
Observations	851,890	871,016	807,858
R-squared	0.007	0.008	0.008

Control variables (estimates not displayed) are age and age squared of the tax filer, total number of dependents in the household and its squared term, state fixed effects, and year fixed effects. Robust standard errors, clustered by household, are in parentheses. A third order polynomial of the running variable is used.
*** p<0.01, ** p<0.05,* p<0.1.

Table 4: Robustness Tests (Pooled RD)

	$log(wages)_{t+1}$		$log(wages)_t$		$\Delta log(wages)_t$	$\Delta log(wages)_{t+1}$			
						Placebo	\$30-50K	\$50-70K	\$70-100K
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
T	-0.008**	-0.007	0.001	0.001	0.001	-0.001	-0.018**	-0.008†	0.003
	(0.004)	(0.006)	(0.003)	(0.006)	(0.003)	(0.003)	(0.008)	(0.005)	(0.006)
$log(wages)_{-1}$	NO	YES	NO	YES	NO	NO	NO	NO	NO
Observations	849,473	849,473	845,594	849,473	845,536	843,272	266,078	353,791	231,094
R-squared	0.633	0.015	0.633	0.011	0.003	0.003	0.007	0.008	0.007

Control variables (estimates not displayed) are age and age squared of the tax filer, total number of dependents in the household and its squared term, state fixed effects. Column (6) is a placebo test using matched month pairs within the same year (e.g. January ("treated") vs. February ("control"), March ("treated") vs. April ("control"), etc..). Robust standard errors, clustered by household, are in parentheses.

*** p<0.01, ** p<0.05,* p<0.1, † p<0.11.

Figure 1: Density of Households, by Day of Birth

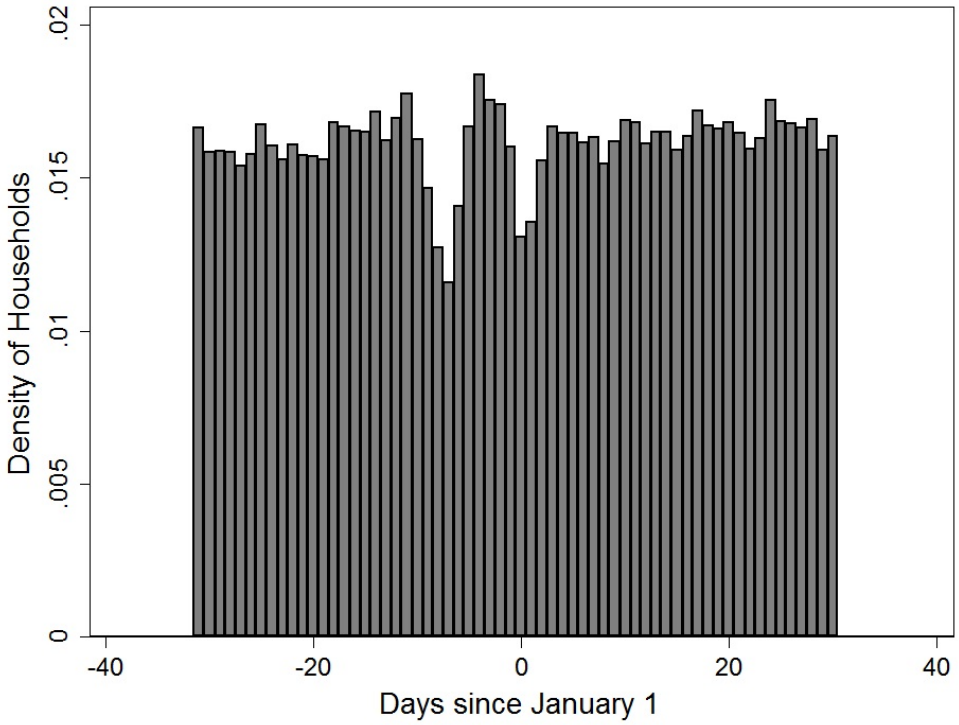

U.S. federal tax filings housed at the IRS's Compliance Data Warehouse, 2004-2011.

Figure 2: CTC Amounts and Log Wages, by Day of Birth

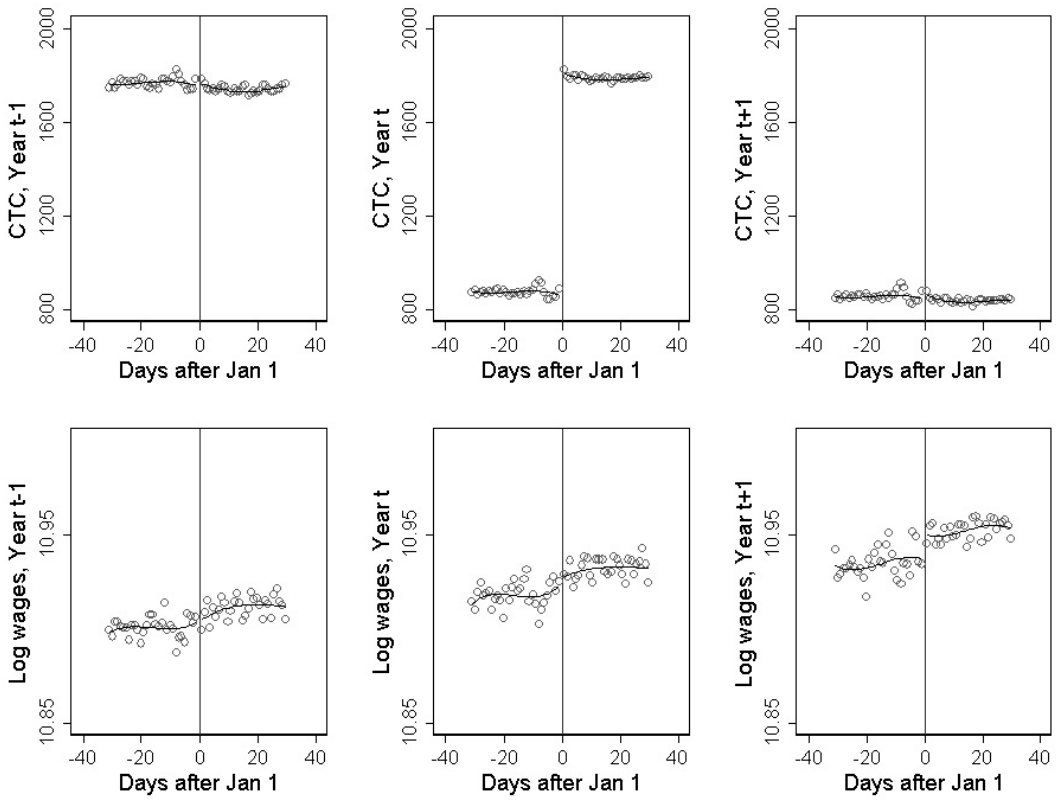

Figures include a third-degree polynomial of the running variable, defined as the number of days after January 1st.

Figure 3: Wage Income Elasticity

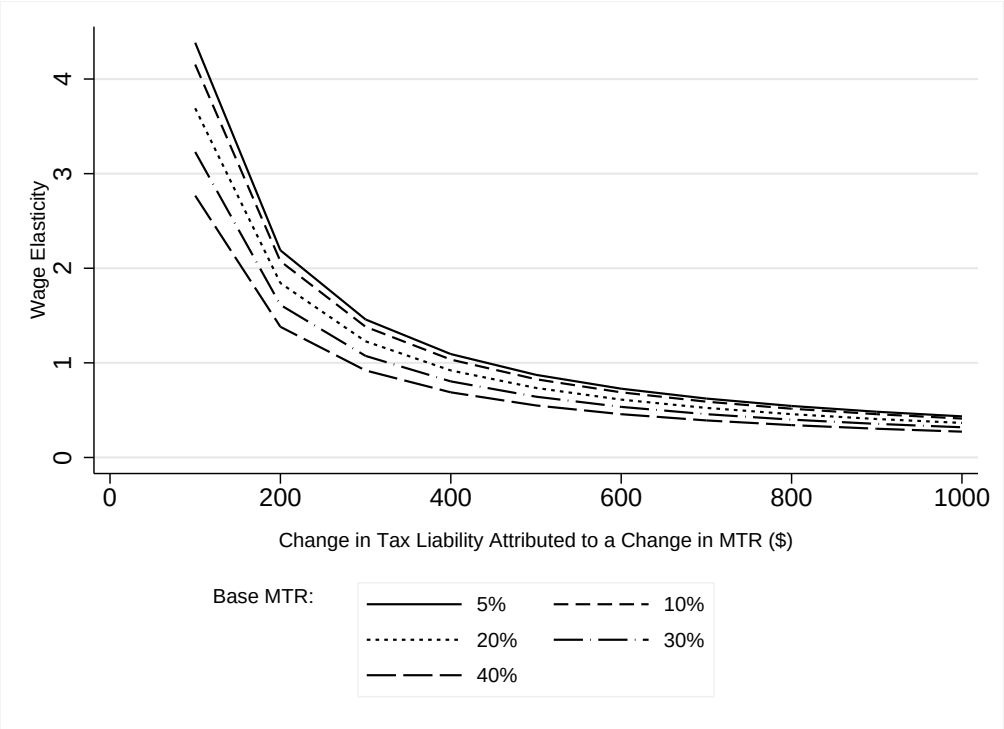

The figure presents the wage income elasticities that result from varying the amount of the increase in tax liability (T) that is due to a perceived increase in the MTR where $T = -D + \tau Y$. \$0 reflects that none of the change is due to the MTR whereas \$1000 reflects that all of the change is due to the MTR. Each of the lines represents a different assumption on baseline household knowledge regarding their MTR. For example, households that believe they face a 5% MTR before realizing the increase in tax liability are represented by the solid black line.

References

Abeler, Johannes and Simon Jäger, "Complex Tax Incentives," *CESifo Working Paper 4231*, 2013.

Altonji, Joseph, "Intertemporal Substitution in Labor Supply: Evidence from Micro Data," *The Journal of Political Economy*, 1986, *93* (3), 176–215.

Auten, Gerald and Robert Carroll, "The Effect of Income Taxes on Household Income," *Review of Economics and Statistics*, 1999, *81* (4), 681–693.

Benjamin, M. Miller and Kevin J. Mumford, "Personal Income Tax Salience: Evidence from the Child and Dependent Care Credit Expansion," *Manuscript*, 2014.

Brown, C. V., "Misconceptions About Income Tax and Incentives," *Scottish Journal of Political Economy*, 1968, *15* (1), 1–21.

Chetty, Raj, "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, 2012, *80 (3)*, 969–1018.

—— **, Adam Looney, and Kory Kroft**, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99*, 1145–77.

Choi, James, David Laibson, and Bridgette Madrian, "Why Does the Law of One Price Fail? An Experiment on Index Mutual Funds," Yale ICF Working Paper No. 08-14 April 2008.

de Bartolome, Charles, "Which Tax Rate Do People Use - Average or Marginal," *Journal of Public Economics*, 1995, *56* (1), 79–96.

Dickert-Conlin, S. and A. Chandra, "Taxes and the Timing of Births," *Journal of Political Economy*, 1999, *107* (1), 161–77.

Dokko, Jane, "The Effect of Taxation on Lifecycle Labor Supply: Results from a Quasi-Experiment," FED Working Paper No. 2008-24 September 2007.

Eissa, Nada, "Taxation and Labor Supply of Married Women: The Tax Reform Act of 1986 as a Natural Experiment," NBER Working Paper 5023, February 1995.

—— **and Jeffrey B. Liebman**, "Labor Supply Response to the Earned Income Tax Credit," *Quarterly Journal of Economics*, 1996, *111* (2), 605–637.

Elisson, Glenn, "Bounded Rationality in Industrial Organization," in R. Blundell, W. Newey, and T. Persson, eds., *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Vol. 2, Cambridge University Press, 2006, pp. 142–174.

Feenberg, Dan and Elizabeth Coutts, "An Introduction to the Taxsim Model," *Journal of Policy Analysis and Management*, 1993, *12* (1), 189–194.

Feldman, Naomi and Bradley Ruffle, "The Impact of Tax Exclusive and Inclusive Prices on Demand," *American Economic Journal: Economic Policy*, 2014, *Forthcoming.*

___ , Jacob Goldin, and Tatiana Homonoff, "Raising the Stakes..," *Working Paper*, 2014.

Feldstein, Martin, "The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act," *Journal of Political Economy*, 1995, *103* (3), 551–572.

Finkelstein, Amy, "E-ZTax: Tax Salience and Tax Rates," *Quarterly Journal of Economics*, 2009, *124* (3).

Fujii, E. T. and C. B. Hawley, "On the Accuracy of Tax Perceptions," *Review of Economics and Statistics*, 1988, *70* (2), 344–347.

Gabaix, Xavier and David Laibson, "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *Quarterly Journal of Economics*, 2006, *121* (2), 505–540.

Goldin, Jacob, "Optimal Tax Salience," Working Paper, 2013.

Goolsbee, Austan, "What Happens When You Tax the Rich? Evidence from Executive Compensation," *Journal of Political Economy*, 2000, *108* (2), 352–378.

Gruber, Jonathan and Emmanuel Saez, "The Elasticity of Taxable Income: Evidence and Implications," *Journal of Public Economics*, 2002, *84* (1), 1–32.

Heidhues, Paul, Botond Kőszegi, and Takeshi Murooka, "Inferior Products and Profitable Deception," *Working Paper*, 2014.

Homonoff, Tatiana and Jacob Goldin, "Smoke Gets in Your Eyes: Cigarette Tax Salience and Regressivity," *American Economic Journal: Economic Policy,*, 2013, *5* (1), 302–36.

Hossain, Tanjim and John Morgan, "...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay," *B.E. Journals in Economic Analysis and Policy: Advances in Economic Analysis and Policy*, 2006, *6* (2), 1–27.

Kopczuk, Wojciech, "The Optimal Elasticity of Taxable Income," *Journal of Public Economics*, 2005, *84* (1), 91–112.

LaLumia, Sara, James Sallee, and Nicholas Turner, "New Evidence on Taxes and the Timing of Birth," *American Economic Journal: Economic Policy*, 2014. Forthcoming.

Liebman, Jeffrey B. and Richard J. Zeckhauser, "Schmeduling," Manuscript, Harvard, October 2004.

Looney, Adam and Monica Singhal, "The Effect of Anticipated Tax Changes on Intertemporal Labor Supply and the Realization of Taxable Income," FED Working Paper No. 2005-44 September 2006.

McClelland, Robert and Shannon Mok, "A Review of Recent Research on Labor Supply Elasticities," CBO Working Paper No. 2012-12 October 2012.

Peterman, William, "Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity," *Working Paper*, 2013.

Poterba, James M. and Andrew Samwick, "Taxation and Household Portfolio Composition: US Evidence from the 1980s and 1990s," *Journal of Public Economics*, 2003, *87* (1), 5–38.

Reck, Daniel, "Taxes and Mistakes What's in a Sufficient Statistic?," Manuscript 2013.

Romich, Jennifer and Thomas Weisner, "How Families View and Use the EITC: Advance Payment versus Lump Sum Delivery," *National Tax Journal*, 2000, *53* (4), 1245–1265.

Saez, Emmanuel, "Reported Incomes and Marginal Tax Rates, 1960-2000: Evidence and Policy Implications," in James M. Poterba, ed., *Tax Policy and the Economy*, Vol. 18, The MIT Press, 2004, pp. 117–173.

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

Slemrod, J., "Cheating Ourselves: The Economics of Tax Evasion," *Journal of Economic Perspectives*, 2007, *21* (1), 25–48.

Souleles, Nicholas S., "The Response of Household Consumption to Income Tax Refunds," *American Economic Review*, 1999, *89* (4), 947–958.