

Estimating Taxable Income Responses with Elasticity Heterogeneity

Anil Kumar and Che-Yuan Liang

Federal Reserve Bank of Dallas
Research Department
Working Paper 1611

<https://doi.org/10.24149/wp1611r1>

Estimating Taxable Income Responses with Elasticity Heterogeneity*

Anil Kumar[&]
Che-Yuan Liang[#]

November 1, 2016
Revised: March 29, 2017

Abstract: We explore the implications of heterogeneity in the elasticity of taxable income (ETI) for tax-reform based estimation methods. We theoretically show that existing methods yield elasticities that are biased and lack policy relevance. We illustrate the empirical importance of our theoretical analysis using the NBER tax panel for 1979-1990. We show that elasticity heterogeneity is the main explanation for large differences between estimates in the previous literature. Our preferred, newly suggested method yields elasticity estimates of approximately 0.7 for taxable income and 0.2 for broad income.

Keywords: elasticity of taxable income, elasticity heterogeneity, tax reforms, panel data, preference heterogeneity

JEL classification: D11, H24, J22

* We thank Lennart Flood, Alexander Gelber, Seth Giertz, Wojciech Kopczuk, Erik Lindqvist, Matthew Rutledge, Håkan Selin, Caroline Weber, seminar participants at Uppsala Center for Fiscal Studies and Department of Economics at Uppsala University, the Workshop on Public Economics and Public Policy in Copenhagen, the 28th Annual Conference of the EEA in Gothenburg, the 2014 AEA Meeting in Philadelphia, the 15th Journées Louis-André Gérard-Varet in Aix-en-Provence, the 4th SOLE/EALE World Conference in Montréal, the Workshop on Behavioral Responses to Income Taxation in Mannheim, the 91st Annual Conference of the Western Economic Association International in Portland, the Federal Reserve System Applied Microeconomics in Cleveland, and the 72nd Annual Congress of the International Institute of Public Finance at Lake Tahoe for valuable comments and suggestions. We are also grateful to Michael Weiss for generous help with the manuscript. A previous version of the paper has also circulated as “The Taxable Income Elasticity: A Structural Differencing Approach”. The Jan Wallander and Tom Hedelius Foundation, the Swedish Research Council for Health, Working Life and Welfare (FORTE), and the Uppsala Center for Fiscal Studies (UCFS) are acknowledged for their financial support. The views expressed here are those of the authors and do not necessarily reflect those of Uppsala University, the Federal Reserve Bank of Dallas, or the Federal Reserve System.

[&] Research Department, Federal Reserve Bank of Dallas; e-mail: anil.kumar@dal.frb.org

[#] Institute of Housing and Urban Research, Uppsala Center for Fiscal Studies, and Department of Economics, Uppsala University; e-mail: che-yuan.liang@nek.uu.se, corresponding author

1. Introduction

The responsiveness of taxable income to tax-rate changes is a widely recognized and important public finance research question. Following the seminal work of Feldstein (1995, 1998), a large body of literature has emerged regarding estimation of the elasticity of taxable income (ETI) with respect to the marginal net-of-tax rate¹ at the observed income level. This literature has generated a wide range of estimates with vastly different implications for optimal tax policy. With significant tax reform well within sight after the recent US elections, evaluating and interpreting the policy consequences of these estimates has assumed particular importance. Estimates obtained using different methods, even for the same reform, remain strikingly different. As an example, previous research on the impact of the tax cuts in the Tax Reform Act of 1986 (TRA86) has produced ETI estimates ranging from 0.2 to 3 (e.g., Feldstein, 1995; Auten and Carroll, 1999; Mofitt and Wilhelm, 2000; Gruber and Saez, 2002; Kopczuk, 2005; Weber; 2014).

The previous literature primarily used instrumental variable regression of the change in the log of taxable income on the change in the log of observed marginal net-of-tax rate. Instruments are required because the observed tax rate is mechanically a function of income; therefore, the change in observed net-of-tax rate is endogenous to the change in taxable income. The most widely used instrument is the net-of-tax rate change constructed holding real taxable income fixed at the base-year income level prior to the tax change. As mentioned by Gruber and Saez (2002), instruments exploiting variation in tax-rate changes due to tax reform are invalid if they are correlated with unobserved trend heterogeneity in income changes. If income is mean reverting, unobserved year-to-year variation in income can cause a positive trend heterogeneity bias. Furthermore, unobserved shocks to taxable income due to widening income distribution – driven by such factors as trade or technological change – can cause negative trend heterogeneity bias. While the search for valid instruments addressed concerns due to trend heterogeneity (e.g., Kopczuk, 2005; Blomquist and Selin, 2010; Weber, 2014; Burns and Ziliak, 2016), previous research has ignored the consequences of different elasticities among individuals owing to, e.g., skill differences. Such heterogeneity is typically an essential component of many models in the theoretical optimal taxation literature (e.g., Mirlees, 1971).

In this paper, we introduce elasticity heterogeneity in the estimation of the ETI in the standard IV setting in first-differences² and make four contributions. First, we show that elasticity heterogeneity, in addition to trend heterogeneity, is an important source of bias. Instruments used in the literature are invalid because they are by construction endogenous to elasticity heterogeneity. Second, we show that different instruments attempt to estimate weighted averages of individual elasticities with *different* weighting functions. *None* of these weighted averages is policy relevant. Third, we propose potentially valid instruments for estimating more policy relevant weighted-average ETIs. Finally, we illustrate the importance of elasticity heterogeneity using the NBER tax panel for 1979-1990 and present new policy relevant estimates after disentangling and quantifying the various sources of bias. We show

¹ The net-of-tax rate is one minus the tax rate. See Saez et al. (2012) for a review of the literature.

² Blomquist et al. (2014) developed a non-parametric method that allows *general* heterogeneity. However, their setting in levels does not nest the standard setting.

that accounting for elasticity heterogeneity helps reconcile the wide variation in ETI estimates arising from the methods in Feldstein (1995), Gruber and Saez (2002), Saez et al. (2012), Weber (2014), and Burns and Ziliak (2016).

The intuition behind the elasticity heterogeneity bias can be illustrated using the estimated impact for TRA86 presented in Table 2 of Feldstein (1995). While the treated group (those with highest pre-reform income and marginal tax rate) received a marginal net-of-tax rate increase of 42% in the post-reform period, the increase for the control group (those with somewhat lower pre-reform income and marginal tax rate) was just 25%. The difference in taxable income change (treated *minus* control) of 51% divided by the difference in net-of-tax rate change of 17% yielded the implied ETI estimate of 3. As noted by Navratil (1995) and Saez et al. (2012, p.26), when the control group also faces a tax change, Feldstein's grouping method is consistent only if the two groups have identical elasticities. However, individuals with different base-year income, *ceteris paribus*, have different elasticities.

Subsequent panel studies did not use grouping methods and, instead, exploited the entire continuous variation in the net-of-tax rate change as base-year income varies. Gruber and Saez (2002) suggested pooling several first-differences to exploit base-year *income-by-year* variation, which allows addressing *trend* heterogeneity by controlling for base-year income. Weber (2014) and Blomquist and Selin (2010) argued that replacing base-year income with lagged base-year income and mid-year income, respectively, would better account for trend heterogeneity bias. We show that the identifying income-by-year variation is endogenous to *elasticity* heterogeneity also for these methods.

While tax-rate changes vary across the income distribution and year, they also vary *within any given income level and year* in a way that depends on demographic factors, such as state of residence, filing status, and number of children. The literature did not exploit such variation, possibly thinking it appeared insufficient. Contrary to conventional belief, however, using the NBER-TAXSIM model, we show that tax-rate changes due to TRA86 vary substantially even at given income levels and years.

We identify two types of potentially valid instruments. First, we propose using income-by-year residualized instruments that remove endogenous income-by-year variation in tax-rate changes from previous instruments. Second, we argue that tax-rate changes at *constant* income levels are uncorrelated with income-by-year variation. The first-dollar tax-rate change, e.g., is such an instrument, and its level version has been widely used in the literature on estimating tax price impact on charitable contributions, 401(k) contributions, capital gains realization, and labor supply.³

An important motivation for our proposed instruments is that tax reform typically changes entire tax schedules, involving multiple tax brackets. Individuals in a particular bracket may potentially react not only to the tax-rate change in that bracket but also to tax-rate changes in other brackets if they switch brackets. This raises the issue about who contributes to identification of an estimated elasticity when using different instruments. We show that while valid instruments yield consistent weighted averages of individual elasticities, the weighting function differs across instruments. Similar in spirit to the local average treatment effect (LATE) in the treatment effects literature, the instruments yield local (weighted

³ See, e.g., Triest (1998) and Keane (2011) for reviews.

average) ETIs where the weight given to each elasticity depends on how strongly the observed net-of-tax rate change is correlated with the instrument. A first-dollar tax-rate change, e.g., most strongly affects the observed tax-rate change of low-income individuals.

We prove that compared with other instruments, the base-year net-of-tax rate change gives the greatest weight to relatively inelastic individuals, as more elastic individuals have a greater likelihood of moving between brackets in response to a base-year tax-rate change. Therefore, the observed tax rates of more elastic individuals are relatively more responsive to tax-rate changes in brackets other than their base-year bracket. In particular, a completely inelastic individual never switches bracket and has an observed tax-rate change equal to the base-year tax-rate change.

The instruments discussed use only a small part of the variation in tax-rate change across the income distribution during a tax reform, which affects precision. Furthermore, the local ETIs are not policy relevant because they only partially capture the effects of the collection of tax-rate changes in the data. One way to account for effects of changes in the entire tax structure is to use multiple constant-income net-of-tax rate change instruments. We propose constructing a single synthetic average net-of-tax rate change instrument that is a weighted average of net-of-tax rate changes across the entire income distribution. We show that weighting each constant-income net-of-tax rate change by the income level's observed probability density yields an ETI analogous to the average treatment effect on the treated (ATT) in the treatment effects literature.

The standard ETI, estimated in much of the literature, is limited because it is measured with respect to the *observed* net-of-tax rate, even as tax structure change is expressed in terms of a set of tax-rate changes at different income levels. With progressive tax rates, a 1% tax reduction at each income level would lower observed tax rates by less than 1%, as some individuals respond to the tax reduction by increasing their income and moving to a higher tax bracket. From a policy perspective, the reduced-form estimates of valid instruments represent *policy* elasticities (similarly defined as in Hendren, 2016), measuring income responses with respect to mechanical tax-rate variables under policy control. Our (weighted) average net-of-tax rate change instrument yields a policy elasticity that is more informative for efficiency analysis than the standard ETI. Specifically, our methodology accounts for the nonlinear budget set complications discussed by Blomquist and Simula (2016).

Our primary empirical finding is that the average net-of-tax rate change yields an ETI of around 0.7. The estimate is robust to inclusion of income control functions and demographic controls and even to inclusion of year-specific versions of these covariates. Furthermore, it is relatively insensitive to using only demographic group-level variation in tax-rate changes for identification. We argue that these results provide evidence for instrument validity.

Our instrument also yields a reduced-form taxable income policy elasticity of 0.46, which is around 70% of the IV estimate. This implies that changing the tax structure by an amount that increases observed net-of-tax rates at the base-year income level by 1% increases taxable income by 0.46%. Furthermore, our instrument yields a broad income elasticity of 0.21 and a broad income reduced-form policy elasticity of 0.13.

We also estimate an ETI of 0.26 using the base-year net-of-tax rate change instrument proposed in Gruber and Saez (2002). We then isolate the continuous base-year income-by-year variation, which is similar in spirit to the variation used by the grouping methods in

Feldstein (1995) and Saez et al. (2012). This method yields estimates of 1.0 to 1.3. On the other hand, using the base-year income-by-year residualized variation yields a consistent ETI of around 0.2. The discrepancy between the estimates of 0.2 and 1.0 to 1.3 reflects a large positive elasticity heterogeneity bias. On the other hand, the discrepancy between the consistent base-year ETI of 0.2 and the ETI of 0.7 for our newly suggested average net-of-tax change instrument reflects that the base-year instrument significantly overweights low-elasticity individuals. We also reproduce an ETI estimate of 0.50 for a Weber (2014)-type net-of-tax rate change instrument evaluated at base-year income lagged two years.

Saez et al. (2012, p.28) offered two explanations for divergence across estimates in the literature. First, they argued that using continuous instruments capturing minor individual-level tax-rate changes leads to lower estimates because individuals are less likely to respond to such rate changes. Second, they claimed that trend heterogeneity could account for much of the sensitivity in estimates across various methods. We find compelling evidence of alternative explanations. We show that the grouping estimates (1 to 3 in, e.g., Feldstein, 1995) were larger than the subsequent ungrouped estimates (0.2 to 1.5 in, e.g., Gruber and Saez, 2002; Weber, 2014) mainly because grouping methods exclude tax-rate variation within given income levels and years, and therefore, suffer from a larger elasticity heterogeneity bias. We also show that the discrepancies between the ungrouped estimates are primarily due to differences in how each elasticity is weighted.

2. Theoretical framework

2.1 Basic model in levels

The taxable income decision problem is such that the individual chooses (Y, c) to maximize utility $u(Y, c)$ subject to a budget constraint $c(Y)$ and $Y \geq 0$, where Y is gross taxable income and c is consumption. The set of points $\{Y, c(Y): Y \geq 0\}$ is the budget frontier of the budget set $\{(Y, c): 0 \leq Y, 0 \leq c \leq c(Y)\}$. We work with a standard iso-elastic quasi-linear utility function with two parameters:

$$u(c, Y; \beta_i, \alpha_i) = c - \frac{\exp(\alpha_i)^{-\frac{1}{\beta_i}}}{1 + \frac{1}{\beta_i}} Y^{1 + \frac{1}{\beta_i}}. \quad (1)$$

$e = (\beta_i, \alpha_i)$ are preference parameters, and subscript i indexes individuals. With locally nonsatiated preferences, the individual consumes all its net income in our static model.

The budget constraint depends on the tax (and transfer) system according to:

$$c(Y) = Y - T(Y) + c_0, \quad (2)$$

where $T(Y)$ expresses net taxes as a function of gross taxable income and where c_0 is net income from sources other than taxable income. We assume that $T(\cdot)$ is exogenous to c_0 . Without loss of generality, for a continuously differentiable budget constraint, the tax schedule/structure can be described by the marginal net-of-tax rate function $t(Y) = dc(Y)/dY = -dT(Y)/dc$. We work with the natural logarithms of Y and t :

$$y = \ln Y, \quad \tau(y) = \ln t(y). \quad (3)$$

The set of tax-rate parameters $\boldsymbol{\tau} = \{(\tau(y): y \geq 0, c_0)\}$ is an alternative way to fully characterize the shape of budget constraint/frontier/set. Because the government sets the tax schedule by setting the tax rate at each income level, e.g., the first-dollar tax rate, the second dollar tax rate, etc., $\boldsymbol{\tau}$ are tax policy variables (allowed to be individual-specific).⁴

An optimal choice of the observed (log of gross taxable) income y^* is given by the first-order condition of an optimization problem with the convex preferences in Eq. (1) if the budget constraint in Eq. (2) is concave, i.e., if marginal tax rates are progressive. Plugging y^* back into $\tau(\cdot)$ yields the observed (log of marginal) net-of-tax rate τ^* .⁵ We get the following system of simultaneous equations:

$$y^*(\beta_i, \alpha_i; \boldsymbol{\tau}) = \operatorname{argmax}_y u(y, c(y)) = y(\tau^*; \beta_i, \alpha_i) = \beta_i \tau^* + \alpha_i, \quad (4)$$

$$\tau^*(\beta_i, \alpha_i; \boldsymbol{\tau}) = \tau(y^*; \boldsymbol{\tau}). \quad (5)$$

A consequence of quasi-linear utility is that there is no income effect that depends on c_0 . The Slutsky condition with a positive substitution effect then implies $\beta_i \geq 0$. From the point of view of Eq. (4), $\beta_i = dy^*/d\tau^*$ represents the (both uncompensated and compensated) elasticity of taxable income with respect to the observed net-of-tax rate (ETI), whereas α_i represents potential taxable income without taxes (in which case $\tau^* = 0$).

We introduce unobserved preference heterogeneity through the error terms (b_i, a_i) , and we let β and α be population-average parameters according to:

$$\beta_i = \beta + b_i, \quad \alpha_i = \alpha + a_i, \quad (6)$$

where $E(b_i) = E(a_i) = 0$. Preference heterogeneity captures differences in taste for work and reflects that income differs between individuals with the same budget set. b_i represents heterogeneity in income that is tax-rate dependent, and a_i denotes heterogeneity in income that is tax-rate independent.⁶ While we allow β_i to vary across individuals, we keep the functional-form assumption that it is constant for each individual. We do not make any distributional assumptions on the error terms. Most empirical work on taxable income allowed one-dimensional preference heterogeneity through α_i .⁷ The optimal taxation literature also typically assumes just one source of heterogeneity, but in this case, it is skill or ability heterogeneity that leads to heterogeneity in β_i in our setting (e.g., Mirlees, 1971; Saez, 2001).

⁴ For the individual, Y and c are variables, whereas $\boldsymbol{\tau}_i$, β_i , and α_i are parameters. For the government and in the estimation, $\boldsymbol{\tau}_i$ are variables, and we want to identify some function of β_i and α_i .

⁵ Like the literature using panel data methods, we do not explicitly model location on kink points in piecewise linear budget frontiers leading to a tax function that is not continuously differentiable. The model here can be augmented according to Liang (2014), which would allow using the tax rate from below or above for individuals at kink points, and that would not affect empirical results.

⁶ The estimated coefficient for the net-of-tax variable in a specification that ignores income effects when such effects exist would represent a mixture of substitution and income effects. This mixture would be individual-specific, which we allow, even if substitution and income effects were constant across individuals.

⁷ Empirical work on labor supply using structural nonlinear budget set models in levels often allow several normally distributed error terms, e.g., the Hausman-type of model (Burtless and Hausman, 1978; Hausman, 1995) and the discrete-choice model (Dagsvik, 1994; van Soest, 1995; Hoynes, 1996; Keane and Mofitt, 1998).

Let us start the analysis with the case with linear budget sets in which there is only one net-of-tax rate $\tau(y) = \bar{\tau}$ that is constant for each budget set and $c = \bar{\tau}Y + c_0$. Plugging in the budget constraint in Eq. (5) into the first-order condition in Eq. (4) gives:

$$y^* = \beta_i \bar{\tau} + \alpha_i. \quad (7)$$

y^* is a function of only $\tau^* = \bar{\tau}$. For each individual, $\beta_i = dy^*/d\bar{\tau}$.

Because $E_e(y^*|\bar{\tau}) = \beta\bar{\tau} + \alpha$, $\beta = E_e(\beta_i) = dE_e(y^*|\tau^*)/d\tau^*$ represents the population-average aggregate ETI. Eq. (7) is a random coefficient model (Wald, 1947). Assuming that τ is statistically independent from e , regressing y^* on $\bar{\tau}$ gives $\beta_{OLS} = \sigma_{y^*,\bar{\tau}}/\sigma_{\bar{\tau}}^2 = \beta$, where σ and σ^2 denote the covariance and variance, respectively.

The taxable income literature handles budget set nonlinearities by assuming that individuals behave according to budget sets linearized at observed income levels. A rationale for this procedure is that the optimal choice is the same on the linearized and nonlinear budget sets (Hausman, 1985; Mofitt, 1990). Plugging in a nonlinear budget constraint into the first-order condition gives:

$$y^* = \beta_i \tau^*(\beta_i, \alpha_i; \tau) + \alpha_i, \quad (8)$$

which is a correlated random coefficient model. Hastie and Tibshirani (1993) called it a varying coefficient model with endogenous regressors. Using Eq. (6), we can rewrite Eq. (8) as $y^* = \beta\tau^* + b_i\tau^* + \alpha_i$. The problem of estimating β by regressing y^* on τ^* is that $\tau^*(e; \tau)$ is correlated with the error term $b_i\tau^* + \alpha_i$ as both are functions of e .⁸

The fundamental source of bias is that, for each τ , e is correlated with both y^* and τ^* due to τ^* mechanically being a function of y^* . In Figure 1, we provide an example with two individuals $i = 1, 2$ with different preferences (β_i, α_i) on a budget set with two tax brackets/segments indexed by superscript $s = 1, 2$ with net-of-tax rates τ^s . They choose $y_i^* = y^{s=i}$ and $\tau_i^* = \tau(y_i^*) = \tau^{s=i}$. It cannot be the case that $\alpha_1 = \alpha_2$ and $\beta_1 = \beta_2$,⁹ and τ_i^* is therefore correlated with y_i^* , β_i , and α_i . Clearly, the OLS estimate of y^* on τ^* is negative, and does not yield a consistent estimate of the positive ETI. In general, the OLS estimate contains a negative simultaneity bias.

⁸ This model is similar to the canonical empirical return-to-schooling model in Card (2001), where y^* is earnings, τ^* is schooling, β_i is marginal return to schooling, and α_i is ability. While both schooling and the observed net-of-tax rate are simultaneously determined endogenous outcome variables, a theoretical difference is that earnings do not affect schooling whereas taxable income affects the observed net-of-tax rate. The regressor is endogenous because of a reverse causality problem in our case. Unlike schooling, we also know *all* determinants of the observed net-of tax rate (taxable income and the tax function).

⁹ An implication of elasticity heterogeneity is that only individuals with $\frac{\bar{y}-\alpha_i}{\tau^2} < \beta_i < \frac{\bar{y}-\alpha_i}{\tau^1}$ bunch at the kink point $y^* = \bar{y}$ (applying the condition in Burtless and Hausman, 1978 to our model). The bunching method (Saez et al., 2010; Chetty et al., 2011) can therefore not say anything about the average ETI of individuals with ETI values outside this interval. Further explorations of the bunching method with elasticity heterogeneity are left for future research.




Figure 1. Negative correlation between taxable income and the observed net-of-tax rate

2.2 Introducing panel dimension

With panel data, individual-specific heterogeneity can be differenced away. Let subscript t index years, and drop superscript $*$ for observed variables for notational simplicity. Then:

$$\Delta_{it}y = y_{i,t+1} - y_{it}, \quad \Delta_{it}\tau = \tau_{i,t+1}(y_{i,t+1}) - \tau_{it}(y_{it}), \quad (9)$$

where $\tau_{it}(y_{it}) = \tau(y_{it}; \tau_{it})$ depends on base-year income y_{it} .

We introduce dynamics in the preference error terms in order to capture common panel complications. Without loss of generality, we let $\beta_{it} = \beta_i$ be fixed over time. On the other hand, we allow the α_i to contain a permanent income component α_{it}^p and a transitory income component α_{it}^v . We specify changes in preference parameters and income according to:¹⁰

$$\alpha_{it} = \alpha_{it}^p + \alpha_{it}^v, \quad \Delta_{it}\alpha_{it}^p = g^p(\alpha_{it}^p) + \alpha_{it}^{p\epsilon}, \quad \Delta_{it}\alpha_{it}^v = g^v(\alpha_{it}^v) + \alpha_{it}^{v\epsilon}, \quad (10)$$

$$\Delta_{it}y = \beta_i \Delta_{it}\tau + \Delta_{it}\alpha, \quad \Delta_{it}\alpha = g_{it}^p(\alpha_{it}^p) + g_{it}^v(\alpha_{it}^v) + \alpha_{it}^{p\epsilon} + \alpha_{it}^{v\epsilon}, \quad (11)$$

where $\alpha_{it}^{p\epsilon}$ and $\alpha_{it}^{v\epsilon}$ are error terms with $E(\alpha_{it}^{p\epsilon}) = E(\alpha_{it}^{v\epsilon}) = 0$.

$\Delta_{it}\alpha$ represents an income trend term that can be heterogeneous across y_{it} . g^p could be increasing in α_{it}^p due to widening income distribution, which would lead to permanent income trends $\Delta_{it}\alpha^p$ that are positively correlated with y_{it} . g^v could be decreasing in α_{it}^v due to mean reversion where individuals with high transitory income revert toward lower income levels. That would lead to transitory income trends $\Delta_{it}\alpha^v$ that are negatively correlated with y_{it} .

Estimation of taxable income responses typically starts with Eq. (11), but with constant β_i across individuals. Identification requires tax reforms that lead to differential changes in observed net-of-tax rates across individuals. While some previous models nest our level model (e.g., Blomquist et al., 2014, which allowed multi-dimensional preference heterogeneity), none of them nests our first-difference model.

¹⁰ Our specification encompasses the cases where permanent income grows at a constant rate according to: $\alpha_{i,t+1}^p = \alpha_i^p + g^p + \alpha_{it}^{p\epsilon}$, and where transitory income is serially correlated according to: $\alpha_{i,t+1}^v = g^v \alpha_i^v + \alpha_{it}^{v\epsilon}$, where g^p and g^v are constants.

We can rewrite Eq. (11) as $\Delta y = \beta \Delta \tau + b_i \Delta \tau + g(\alpha_{it}^p, \alpha_{it}^v, \alpha_{it}^{p\epsilon}, \alpha_{it}^{v\epsilon})$. The problem of estimating β by regressing Δy on $\Delta \tau$ is that $\Delta \tau = \Delta \tau(\beta_i, \alpha_{it}^p, \alpha_{it}^v, \alpha_{it}^{p\epsilon}, \alpha_{it}^{v\epsilon}; \Delta \tau_{it}, \tau_{it})$ is correlated with the error term $b_i \Delta \tau + g(\alpha_{it}^p, \alpha_{it}^v, \alpha_{it}^{p\epsilon}, \alpha_{it}^{v\epsilon})$ as both are functions of preference error terms. The first-difference equation is therefore still a correlated random coefficient model. For the simple case without any reform ($\Delta_i \tau = \mathbf{0}$), income trends are positively correlated with Δy , which in turn is negatively correlated with $\Delta \tau$, because some individuals increasing their income switch to tax brackets with higher tax rates. This leads to a first-difference version of the negative simultaneity bias.

2.3 Estimation with instrumental variables

It is well known from Wooldridge (1997) and Heckman and Vytlacil (1998) that estimation with instrumental variables could yield consistent estimates of correlated random coefficient models. In the first-difference setting, let z denote the instrument, let ρ denote the reduced-form estimate, let γ denote the first-stage estimate, and let β_{IV} denote the IV estimate. We can then define and derive the following relationships:

$$\rho = \frac{\sigma_{\Delta y, z}}{\sigma_z^2}, \gamma = \frac{\sigma_{\Delta \tau, z}}{\sigma_z^2}, \beta_{IV} = \frac{\rho}{\gamma} = \frac{\sigma_{\Delta y, z}}{\sigma_{\Delta \tau, z}}, \quad (12)$$

$$\beta_{IV} = \beta^{LATE} + bias^b + bias^a, \quad (13)$$

$$\beta^{LATE} = \frac{\sigma_{E_e(\Delta y | \Delta \tau), E_e(z | \Delta \tau)}}{\sigma_{E_e(\Delta \tau | \Delta \tau), E_e(z | \Delta \tau)}} = \sum_{it} \beta_i w_{it}^{LATE}, \quad (14)$$

$$w_{it}^{LATE} = \frac{\Delta \tau [E_e(z | \Delta \tau) - E_{it}(z)]}{\sum_{it} \Delta \tau [E_e(z | \Delta \tau) - E_{it}(z)]}, \quad (15)$$

$$bias^b = \left[\frac{E_{\Delta \tau}(\sigma_{\beta_i \Delta \tau, z} | \Delta \tau)}{E_{\Delta \tau}(\sigma_{\Delta \tau, z} | \Delta \tau)} - \beta^{LATE} \right] \frac{E_{\Delta \tau}(\sigma_{\Delta \tau, z} | \Delta \tau)}{\sigma_{E_e(\Delta \tau | \Delta \tau), E_e(z | \Delta \tau)} + E_{\Delta \tau}(\sigma_{\Delta \tau, z} | \Delta \tau)}, \quad (16)$$

$$bias^a = \frac{E_{\Delta \tau}(\sigma_{\Delta \alpha_{it, z}} | \Delta \tau)}{\sigma_{\Delta \tau, z}}. \quad (17)$$

The equality in Eq. (13) follows from the law of total covariance. β^{LATE} is the correlation due to variation in budget set changes $\Delta \tau$. Assuming that $\Delta \tau$ and preferences \mathbf{e} are independent, we derive the second equality in Eq. (14) in Appendix A. β^{LATE} represents the exact function of individual elasticities that could be estimated. We refer to any weighted average of individual elasticities as an aggregate ETI. The weight w_{it}^{LATE} depends on the degree of *compliance*, i.e., the correlation between $\Delta \tau$ and z due to $\Delta \tau$. Individuals with a higher correlation contribute more. The ETIs are local in the same sense as the local average treatment effect (LATE) in the treatment effects literature (Imbens and Angrist, 1994; Angrist and Imbens, 1995).

Instrument relevance requires z to be correlated with $\Delta \tau$. The two bias terms $bias^b$ and $bias^a$ reflect correlations due to variation in \mathbf{e} conditional on $\Delta \tau$.¹¹ They are nonzero when z

¹¹ Using the terminology of the treatment effects literature, $\Delta \tau$ measures treatment intensity and z measures treatment intention. Furthermore, β^{LATE} indicates the external validity of β_{IV} , whereas $bias^a$ and $bias^b$ indicate the internal validity of β_{IV} .

is correlated with β_i and $\Delta\alpha_{it}$ for any given $\Delta\tau$. The only way relevance can be achieved without violating the exclusion restriction is by z being correlated with budget set variables and their changes, τ and $\Delta\tau$, which are the only other determinants of $\Delta\tau$ besides \mathbf{e} .¹²

The IV setting in Eqs. (12) to (17) is very general. Using $z = \Delta\tau$ yields the first-difference estimate of Δy on $\Delta\tau$. This estimate is an interesting benchmark because the underlying consistent ETI equals a weighted average elasticity on the taxed β^{ATT} (see Appendix A). This is similar to the weighted average treatment effect on the treated (weighted ATT) that can be estimated in regressions in the treatment effects literature when treatment intensity is continuous. While β^{ATT} is policy relevant unlike most β^{LATE} , the first-difference estimate does not equal it because it yields nonzero bias terms.¹³

w_{it}^{LATE} , and therefore β^{LATE} and β^{ATT} , vary between data sets with different tax reforms producing different collections of budget set changes. β^{LATE} and β^{ATT} are therefore mixtures of preference and budget set parameters and do not represent pure deep universal behavioral parameters that are immutable to the tax system. Slemrod and Kopczuk (2002) demonstrated and explored this insight for the case without elasticity heterogeneity. For a given tax reform, a consistently estimated β^{ATT} accounts for the reform-specific compliance of each individual and is generally the most policy relevant parameter. It is, however, not informative for other types of reforms in terms of predicting behavioral effects. In comparison, β is a deep parameter. However, it only predicts income responses to tax-rate changes conditional on individuals never switching tax brackets, which is only relevant with linear budget sets.¹⁴

Most methods either explicitly used the IV specification in Eq. (12), e.g., Gruber and Saez (2002), or implicitly estimated such specifications, e.g., Feldstein (1995). With the constant elasticity assumption $\beta_{IV} = \beta^{LATE} = \beta^{ATT} = \beta$. This functional form implies $bias^b = 0$ and ignores the elasticity heterogeneity bias, although the literature has widely accounted for the trend heterogeneity bias due to $bias^a$.

Empirical analysis often addressed elasticity heterogeneity by estimating subsample-specific ETIs, sometimes by exploiting variation in tax changes across subsamples (e.g., Kawano et al., 2016). While such methods in some cases can consistently estimate a β^{LATE} (for the full sample), they cannot generally recover β^{ATT} , which requires accounting for the fact that tax-rate changes typically are correlated with elasticity heterogeneity both between and within subsamples.

¹² This is similar to using arguably exogenous institutional characteristics as instruments for schooling in the return-to-schooling application.

¹³ Removing the difference operators in Eqs. (12) to (17) yields an IV in a level setting. Consistency would then require budget sets τ (rather than their changes) to be independent from \mathbf{e} .

¹⁴ Of course, knowing the entire distribution of β_i allows simulating β^{ATT} in *different* tax reforms. Blomquist et al. (2014) showed, however, that pure preference parameters are not generally identified.

3. Estimation with different instruments

3.1 Instruments using income-by-year variation in tax-rate changes

Most instruments in the literature exploit variation in tax-rate changes at *different income levels* due to tax reform. Because individuals have different income, even reforms that lead to the same change in tax schedule for everyone can be exploited. Feldstein (1995) used variation in tax-rate changes across groups based on (pre-reform) base-year income. This grouping method corresponds to using the following instrument:

$$z_0^{\bar{y}}(y_{it}) = 1(y_{it} > \bar{y}), \quad (18)$$

where \bar{y} is the top tax bracket income cutoff.¹⁵ We use subscript 0 to denote base-year income.

In tax reforms, tax-rate changes often vary gradually across multiple tax brackets. To use all the available variation in tax changes, Eq. (18) can be modified by letting $z_0 = c(y_{it})$, where $c(\cdot)$ can be, e.g., a polynomial or a spline. Such an “ungrouped” instrument can assume multiple values and even be continuous.

Base-year instruments may satisfy instrument relevance because $y_{it}(\beta_i, \alpha_{it}^p, \alpha_{it}^v; \tau_{it})$ and $\Delta_{it}\tau = \Delta_{it}\tau(\beta_i, \alpha_{it}^p, \alpha_{it}^v, \alpha_{it}^{p\epsilon}, \alpha_{it}^{v\epsilon}; \Delta\tau_{it}, \tau_{it})$ are correlated as both are functions of $\beta_i, \alpha_{it}^p, \alpha_{it}^v$, and τ_{it} . However, the instruments’ correlation with preference parameters violates the exclusion restriction.¹⁶ While the correlation with permanent and transitory income trends g_{it}^p and g_{it}^v (through α_{it}^p and α_{it}^v) leads to a trend heterogeneity (non-parallel trend) bias, the correlation with β_i leads to an elasticity heterogeneity bias. The reason is that, *ceteris paribus*, individuals with different elasticities have different income, as we saw in Figure 1.¹⁷

In Figure 2, we illustrate a stylized TRA86-example with a budget set with two tax brackets with net-of-tax rates $\tau_t^{s=1,2}$ before the reform and $\tau_{t+Dt}^{s=1,2}$ after the reform. The tax reform results in the net-of-tax changes $\Delta\tau^s = \tau_{t+Dt}^s - \tau_t^s$. There are larger tax reductions at higher income levels with $\Delta\tau^2 > \Delta\tau^1$. Furthermore, there are two individuals with $y_{it} = \beta_{i=1,2}\tau_{it} + \alpha_{i=1,2;t}$ and $\beta_2 > \beta_1$ experiencing income changes, $\Delta y_i = y_{i,t+Dt} - y_{it}$. They locate on tax bracket $i = s$ both before and after the reform, i.e., $\tau_{it}^* = \tau(y_{it}) = \tau_t^{s=i}$.

In this example, no individual switches tax brackets after the reform. For Feldstein’s instrument in Eq. (18), the first stage $\gamma = 1$ as $\Delta\tau^* = \Delta\tau^s$. $\beta_{IV} = \rho = (\Delta y_2 - \Delta y_1)/(\Delta\tau^2 - \Delta\tau^1)$ ¹⁸ is the ratio between the income and observed net-of-tax difference-in-differences (DID). The DIDs compare changes *between* tax brackets where

¹⁵ Tax reforms can also be exploited with repeated cross sections and aggregated time-series. Lindsey (1987), Feenberg and Poterba (1993), Slemrod (1996), and Saez (2004) grouped individuals by their observed incomes. As Saez et al. (2012) noted, changes in group composition over time could be an issue without panel data.

¹⁶ These instruments do, however, account for the correlation between $\Delta\tau$ and $(\alpha_{it}^{p\epsilon}, \alpha_{it}^{v\epsilon})$, unlike the first-difference regression without instruments.

¹⁷ This can also be seen from the first-order condition $y^* = \beta_i\tau^* + \alpha_i$. For the entire equation system, we can show that $dy^*/d\beta_i = \tau^*/[1 - \beta_i\partial\tau(y^*)/\partial y] \neq 0$. The sign and magnitude of bias could be different for other utility functions and depend on the degree of correlation between β_i and α_i . The bias is zero only when α_i is a particular function of β_i and the tax schedule which implies one-dimensional heterogeneity.

¹⁸ Note that without random shocks ($\alpha_{it}^{p\epsilon} = \alpha_{it}^{v\epsilon} = 0$), $\beta_{IV} = \beta_{FD}$. This example therefore also illustrates the problem with the first-difference estimate when there is a tax reform that contributes to the identification.

the second bracket individual is the treated and the first bracket individual is the control. The IV estimate, therefore, relates the difference between the thick horizontal arrows to the difference between the vertical arrows.¹⁹ For clarity, but without loss of generality, assume that the first individual is a representative individual not affected by widening income distribution or mean reversion, unlike the second individual. In this case, we can represent the decomposition of $\Delta y_2 = \beta_2 \Delta \tau^2 + g_2^p + g_2^v$ using the thin arrows in the figure. We have $\beta^{LATE} + bias^b = (\beta_2 \Delta \tau^2 - \beta_1 \Delta \tau^1) / (\Delta \tau^2 - \Delta \tau^1) \geq \beta^{LATE} = \beta_2 \geq \beta_1$. Furthermore, we have $bias^a = (g_2^p + g_2^v) / (\Delta \tau^2 - \Delta \tau^1)$.




Figure 2. Elasticity and trend heterogeneity biases

Auten and Carroll (1999) suggested accounting for trend heterogeneity by controlling for base-year income. With only one first-difference, the base-year income control function will soak up most of the variation in the instrument. Gruber and Saez (2002) proposed pooling several first-differences and using variation in tax-rate changes across base-year income levels and years. Based on this idea, we can generalize Eq. (18) as follows:

$$\mathbf{z}_0^{yt}(y_{it}, \mu_t) = c(y_{it})\boldsymbol{\mu}_t. \quad (19)$$

\mathbf{z}_0^{yt} is a vector-valued function, μ_t represents year-fixed effects, and $\boldsymbol{\mu}_t$ is a vector of year dummies. We use a spline for $c(\cdot)$ in Eq. (19). Because the instruments are year-specific, we can control for base-year income by including a control function $c(y_{it})$ as covariates without destroying identification. We can also control for macro-economic shocks correlated with the timing of reforms by including $\boldsymbol{\mu}_t$ as covariates. Because $c(y_{it}(g_{it}^p + g_{it}^v))\boldsymbol{\mu}_t$ is correlated with $g_{it}^p + g_{it}^v$ through y_{it} , conditioning on y_{it} leads to $bias^a = 0$ in Eq. (17),²⁰ as

¹⁹ The length of arrows is meant to represent the magnitude of the relative income change (rather than the absolute change) and the relative net-of-tax change measured by the rotation (rather than the vertical distance).

²⁰ Weber (2014) correctly argued that the base-year control function estimates are biased estimates of the two separate permanent and temporary income trends. However, she also argued that because of this, the control

$\sigma_{g_{it}^p+g_{it}^v,c(y_{it}(g_{it}^p+g_{it}^v))\mu_t}|c(y_{it}(g_{it}^p+g_{it}^v)) + \mu_t = 0$. Using income-year interactions as instruments while controlling for the non-interacted variables, therefore, overcomes the trend heterogeneity bias.

In the example in Figure 2, instruments based on income-by-year variation can be used if we have an additional cross-section of pre-reform first-differences with the same two individuals experiencing no tax-rate changes. In the pre-reform first-differences, the individuals are also affected by the second and third horizontal trend arrows and $\Delta y_{2,t-Dt} = y_{2t} - y_{2,t-Dt} = g_2^p + g_2^v$. We could therefore eliminate these trends from the reform first-difference.

Using income-by-year variation, however, does *not* address elasticity heterogeneity; the income control function identified from other pre-reform years only captures effects that are general across years, while β_i interacts with $\Delta\tau$ that is year-specific. Formally, $\sigma_{\beta_i\Delta\tau,c(y_{it}(\beta_i))\mu_t}|c(y_{it}(\beta_i)) + \mu_t \neq 0$ and $bias^b \neq 0$ in Eq. (16). Year-specific income control functions can account for the bias but would also remove the identifying variation. In Figure 2, the arrow numbered by 1 does not affect the individuals in the pre-reform first-difference. There is an elasticity heterogeneity bias even with parallel trends or with identical pre-reform first-differences for the two individuals.

In the TRA86-application in Table 2 of Feldstein (1995), the top income group received tax reductions that increased net-of-tax rates by 42%, whereas the subsequent high-income group received tax reductions that increased net-of-tax rates by 25%. This provides a numerical example for Figure 2. The additional income increase of the top income group can not only be attributed to the additional 17% net-of-tax increase but also reflects a differential response to the first common 25% net-of-tax increase.²¹ Removing this differential response between the two groups requires pre-reform first-differences of 25% net-of-tax increase in each group.

For Feldstein's grouping method, Navratil (1995) and Saez et al. (2012) noted that consistency requires either the same elasticity across groups or a control group that remains untreated, i.e., without a tax change. The control group is, however, rarely untreated because tax reform typically introduces a bundle of new programs, some of which affects everybody. Our discussion shows that even variation in tax-rate changes that is continuous across the income distribution or that vary by base-year income *and* year is contaminated by elasticity heterogeneity.²²

function cannot account for trend heterogeneity bias. We believe it can. Our extended example with two pooled first-differences in Figure 2 below illustrates this.

²¹ The intuition is general and applies also to reduced-form tax reform evaluation methods. Eissa and Liebman (1996) provides an example from the labor supply literature. Lone mothers with children were affected by EITC+TRA86, and lone mothers with children were only affected by TRA86. A comparison of the two groups cannot provide the effect of EITC unless both groups responded equally to TRA86.

²² Because controlling for income could alleviate or worsen elasticity heterogeneity bias, we cannot attribute the discrepancy between conditional and unconditional estimates to the trend heterogeneity bias alone. Elasticity heterogeneity also leads to idiosyncratic year-specific non-parallel responses to universal tax reforms in *pre-reform* periods. Such reforms could be subtle, such as implicit tax code revisions due to inflation leading to bracket creep type of effects (Saez, 2003). This would invalidate using income control functions to account for the trend heterogeneity bias.

3.2 Net-of-tax change instruments

Auten and Carroll (1999) and Gruber and Saez (2002) used net-of-tax change constructed holding real taxable income fixed at the base-year income level as an instrument:

$$\Delta\tau_0 = \Delta_{it}\tau(y_{it}; \Delta\tau_{it}) = \tau_{i,t+Dt}(y_{it}) - \tau_{it}(y_{it}). \quad (20)$$

We propose a procedure to remove the endogenous base-year income-by-year variation by regressing $\Delta\tau_0$ on year-specific income functions $c(y_{it})\mu_t$, where we use a local polynomial for $c(\cdot)$. $\Delta\tau_0$ can then be decomposed into:

$$\Delta\tau_0 = c(y_{it})\mu_t + \varepsilon, \quad (21)$$

$$\Delta\tau_0^{yt} = c(y_{it})\mu_t = \Delta_{it}\tau(y_{it}, \mu_t), \quad (22)$$

$$\Delta\tau_0^{-yt} = \varepsilon = \Delta_{it}\tau(y_{it}; \Delta\tau_{it} | y_{it}, \mu_t). \quad (23)$$

The predicted net-of-tax change $\Delta\tau_0^{yt}$ is conceptually the expectation of $\Delta\tau_0$ over observations with the same base-year income in the same year.²³ It is a nonlinear function of base-year income and year similar to the income-year interactions in Eq. (19).

The residualized net-of-tax change $\Delta\tau_0^{-yt}$ captures the remaining variation in tax-rate changes *within* each base-year income level and year. An alternative to using $\Delta\tau_0^{-yt}$ is to use $\Delta\tau_0$ and include $c(y_{it})\mu_t$ as covariates. The residualized variation comes purely from differential tax-schedule changes over years across demographic groups that is uncorrelated with $c(y_{it})\mu_t$. Consistency requires that *conditional* on income-year interactions, demographic status is independent of preferences. The literature did not exploit this conditional variation, possibly thinking it appeared insufficient. Using the NBER-TAXSIM model, we show that $\Delta\tau_0$ varies substantially in TRA86 even at given income levels and years.

In Figure 3, we illustrate the differences between $\Delta\tau_0$, $\Delta\tau_0^{yt}$, and $\Delta\tau_0^{-yt}$. There is one pre-reform budget set and two post-reform budget sets $k = A, B$. Each budget set contains two tax brackets, where net-of-tax rates are $\tau_t^{s=1,2}$ before the reform and $\tau_{k,t+Dt}^{s=1,2}$ after the reform. Two types of individuals, with $(\beta_{i=1,2}, \alpha_{i=1,2;t})$ are each observed twice on each bracket in the pre-reform budget set and once on each bracket in each of the post-reform budget sets. There are eight observations with y_{it} before the reform and $y_{i,k,t+Dt}$ after the reform, generating four first-differences $\Delta_{ik}y = y_{i,k,t+Dt} - y_{it}$ indicated by the arrows in the figure. For clarity, individuals of type i are observed on bracket $s = i$ both before and after the reform, with $\tau_{it} = \tau(y_{it}) = \tau_t^{s=i}$ and $\tau_{i,k,t+Dt} = \tau(y_{i,k,t+Dt}) = \tau_{k,t+Dt}^{s=i}$, generating $\Delta_{ik}\tau = \Delta_k\tau^{s=i} = \tau_{k,t+Dt}^{s=i} - \tau_t^{s=i}$.

²³ Another way to implement Eq. (22) is to group observations into multiple income groups and assign the group-average income to each observation. To exploit the entire possibly continuous variation in tax-rate changes, our strategy that lets $c(\cdot)$ be a local polynomial corresponds to assigning a synthetic average within an income band to each observation. Note that we reserve the use of “predicted net-of-tax change” for $\Delta\tau_0^{yt}$ although all net-of-tax change instruments are predicted (unlike $\Delta\tau$), and some authors use it for $\Delta\tau_0$.




Figure 3. Income-by-year and residualized variation in tax-rate changes

In this example, $\Delta\tau_0 = \Delta_k\tau^s$, $\Delta\tau_0^{yt} = 0.5(\Delta_A\tau^s + \Delta_B\tau^s)$, and $\Delta\tau_0^{-yt} = 0.5(\Delta_A\tau^s - \Delta_B\tau^s)$. $\Delta\tau_0^{yt}$ groups individuals by brackets (s). It compares individuals at different brackets receiving different average slope rotations, i.e., the changes $\Delta_{1A}y$ and $\Delta_{1B}y$ with $\Delta_{2A}y$ and $\Delta_{2B}y$. $\Delta\tau_0^{-yt}$ groups individuals by tax-schedule changes (k). It compares individuals on the same bracket receiving different slope rotations, i.e., the changes $\Delta_{1A}y$ with $\Delta_{1B}y$ and $\Delta_{2A}y$ with $\Delta_{2B}y$. $\Delta\tau_0^{-yt}$ yields an ETI that is a weighted average of the horizontal difference over the rotational difference between the thin and thick arrows.

Several studies in the literature suggested constructing instruments that are related to $\Delta\tau_0$ by replacing y_{it} with other instrument income y^z . Weber (2014) believed that trend heterogeneity due to mean reversion is not satisfactorily addressed by controlling for base-year income as suggested by Gruber and Saez (2002). She then showed that constructing net-of-tax change instruments based on lags of base-year income $y_{i,t-l}$ mitigates this concern in the limit as l increases, as $y_{i,t-l}$ becomes independent of temporary income. In our application, we use the Weber-type instrument where $y^z = y_{i,t-2}$:

$$\Delta\tau_{-2} = \Delta_{it}\tau(y_{i,t-2}; \Delta\tau_{it}) = \tau_{i,t+Dt}(y_{i,t-2}) - \tau_{it}(y_{i,t-2}). \quad (24)$$

To account for widening income distribution, Weber included a spline in lagged base-year income as covariates, in our case, $c(y_{i,t-2})$, as a proxy for permanent income trends. Blomquist and Selin (2010) made similar remarks about mean reversion and suggested using mid-year income as instrument income. Even these instrument income alternatives are, however, endogenous to elasticity heterogeneity. In the simple example in Figure 2, it is entirely possible that individuals never switch tax brackets. Grouping by lagged and mid-year income would then yield identical estimates as grouping by base-year income.

Demographic variables are correlated with preferences to a much lesser degree than income is, so variation in tax-rate changes by demographics and year is plausibly much cleaner. Including demographic covariates can account for any remaining *trend* heterogeneity

bias. *Year-specific* demographic covariates are, however, needed to account for potential *elasticity* heterogeneity bias.²⁴ Excluded interaction terms between demographic variables still help identification and are likely exogenous. Even when a tax reform appears to be universal, random variation often exists once the entire tax system, including tax credits and deductions, has been accounted for in the budget sets.²⁵ In our application, we explore the inclusion of general and year-specific dummies based on state of residence, marital status, and number of children as covariates.

We also investigate the scope of variation in tax-rate changes by our demographic variables in detail. In particular, we group instruments by our demographic variables and the double and triple interactions between them for each year separately, while including the non-interacted variables as covariates.

Several grouping methods in the labor supply literature exploit variation in tax-rate changes across demographic characteristics. In the EITC-application in Eissa and Liebman (1996), grouping is based on single mothers with or without children. In the labor supply application in Blundell et al. (1998), grouping is based on cohort-education interactions, and they include the non-interacted variables as covariates. Burns and Ziliak (2016) provided a recent taxable income application that groups the base-year net-of-tax change instrument by state-cohort-education interactions, and they include the non-interacted variables as covariates. For these methods to yield consistent estimates, the identifying group-level variation in tax-rate changes must be uncorrelated with income-year interactions. Ensuring parallel trends is not enough. Including covariates is a good remedy, but they need to be year-specific (in a first-difference equivalent setting) to account for elasticity heterogeneity bias.²⁶

3.3 Instruments using variation within income levels and years

By plugging $\Delta\tau_0$ and $\Delta\tau_0^{-yt}$ in Eq. (20) and (23) into Eqs. (12) to (17), we show that $\beta_0^{IV} = \beta_0^{LATE} + bias^b + bias^a$ and $\beta_0^{IV,-yt} = \beta_0^{LATE,-yt}$. Furthermore, the two instruments have the same consistent local ETI, i.e.:

$$\beta_0^{LATE} = \beta_0^{LATE,-yt}. \quad (25)$$

We can, therefore, quantify the local ETI and elasticity heterogeneity bias of $\Delta\tau_0$ using $\Delta\tau_0^{-yt}$.

Instead of removing endogenous income-by-year variation from an invalid instrument such as $\Delta\tau_0$, tax-rate changes within income levels (and years) can be isolated by using net-of-tax changes at *constant* income levels as instruments:

$$\Delta\tau_{\bar{y}} = \Delta_{it}\tau(\bar{y}; \Delta\tau_{i,t}) = \tau_{i,t+Dt}(\bar{y}) - \tau_{i,t}(\bar{y}), \quad (26)$$

where \bar{y} is an income level that is constant across individuals. The first-dollar net-of-tax change is an example of an instrument. Its level version has been widely used in the literature

²⁴ This conclusion rests on the same type of argument used to motivate the need of year-specific income control functions to account for the elasticity heterogeneity bias discussed in Subsection 3.2.

²⁵ This is similar to how the same universal tax code (such as the federal tax rates) often have different effects on tax rates once the entire tax system is accounted for. That type of level (rather than our first-difference) variation in tax rates is used for identification in structural nonlinear budget set methods (e.g., in the discrete-choice method in Dagsvik, 1994; van Soest, 1995; Hoynes, 1996; Keane and Mofitt, 1998).

²⁶ In contrast, $\Delta\tau_0^{-yt}$ is by construction uncorrelated with income-year interactions.

on estimating tax price impact on charitable contributions, 401(k) contributions, capital gains realization, and labor supply.

An important motivation for our proposed instruments is that tax reform typically changes entire tax schedules involving multiple tax brackets. Individuals may potentially react to tax-rate changes across the income distribution. Consider an individual that increases income in response to a base-year tax-rate change and switches to a new adjacent tax bracket. Such an individual may then respond to an adjacent tax-rate change that would affect the observed net-of-tax change. In other words, the individual partially complies with the adjacent net-of-tax change.

While all instruments using variation within income levels and years potentially yield consistent ETI estimates, the contribution of individual elasticities to these estimates differs across instruments. Each instrument yields a local ETI where the weight given to each elasticity depends on the degree of compliance with the instrument. Varying the first-dollar tax change is, e.g., more likely to lead to variation in the observed tax change of individuals with low income compared to individuals with high income. The first-dollar instrument therefore gives more weight to these individuals, and they have, on average, elasticities that are different from other individuals.

Like $\Delta\tau_{\bar{y}}$, $\Delta\tau_0$ only uses a single tax-rate change for each individual, although this change is evaluated at different income levels for different individuals. For this reason, it is tempting to believe that the local (consistent) ETI of base-year instruments ($\Delta\tau_0$ and $\Delta\tau_0^{-y^t}$) is a weighted average of the local ETI of different $\Delta\tau_{\bar{y}}$. In Appendix A, we prove this is not the case. Instead, among all instruments using variation within income levels and years, the base-year instruments minimize the local ETI, i.e.:

$$\underset{y^z}{\operatorname{argmin}} \beta^{LATE} [z = \tau_{i,t+Dt}(y^z) - \tau_{it}(y^z)] = y_{it}. \quad (27)$$

The reason for this result is that the degree of compliance is the most negatively correlated with elasticity heterogeneity for the base-year instruments, resulting in low-elasticity individuals receiving the highest relative weight. This “overweighting” decreases as y^z moves further away from base-year income. The local ETI for net-of-tax change instruments based on, e.g., lagged income is therefore higher than that for the base-year instruments.

In Figure 4, we provide an example of the degree of compliance to different tax-rate changes. Like the example in Figure 3, there is one pre-reform budget set and two post-reform budget sets $k = A, B$, each containing two tax brackets $s = 1, 2$. The second bracket differs between the two post-reform budget sets. There is one individual of type $i = 1$ and two individuals of type $i = 2$. Everyone has the same pre-reform base-year income level y_{it} on the first bracket. Assume for clarity that preferences are time-fixed. After the reform, the first type of individual stays on the first bracket and moves to $y_{1,t+x}$, but the second type of individual switches to the second bracket and moves to $y_{2,k,t+x}$, giving the three first-differences indicated by the arrows in the figure.




Figure 4. Compliance to different tax-rate changes

The non-switcher complies fully with the first bracket base-year net-of-tax change, as $\Delta_1\tau = \Delta\tau^1$ (the dashed arrow).²⁷ If this change varies, observed net-of-tax change would adjust by the same amount. However, the non-switcher does not comply with the adjacent second bracket net-of-tax change – whether this individual faces $\Delta_A\tau^2$ or $\Delta_B\tau^2$ (the thick or thin arrows) does not affect the observed net-of-tax change. On the other hand, the switchers do not fully comply with the base-year net-of-tax change. However, they partially comply with the adjacent net-of-tax change, as the observed net-of-tax change is larger for the switcher receiving the larger second bracket net-of-tax change, i.e., $\Delta_{2B}\tau > \Delta_{2A}\tau$ and $\Delta_B\tau^2 > \Delta_A\tau^2$. In relative terms, compared to non-switchers, switchers comply more with the adjacent net-of-tax change than the base-year net-of-tax change.

Now, the switchers receive a larger income increase but a smaller observed net-of-tax increase compared with the non-switcher. This is only possible if the switchers have a higher elasticity. In general, a more elastic individual responds more to a base-year net-of-tax change and is therefore more likely to be a switcher. In contrast, a completely inelastic individual never switches tax bracket. In other words, high-elasticity individuals are more likely to be switchers and comply relatively less with the base-year net-of-tax change that, on average, gives the greatest weight to low-elasticity individuals.²⁸

3.4 Average net-of-tax change instrument

Each of the instruments discussed so far uses only a small part of the tax-reform variation in tax-rate changes across the income distribution. Discarding the remaining useful variation affects precision. What is more problematic is that each of the local ETIs only captures some of the reform effects and is, therefore, not policy relevant, even for the set of tax-rate changes seen in the data. One way to capture effects of changes in the entire tax structure is to use

²⁷ For individuals that locate in the second bracket already before the reform, the second-bracket net-of-tax change is the base-year net-of-tax change.

²⁸ As long as the first-stage effect of $\Delta\tau_0$ on $\Delta\tau$ is different from one, which is almost always the case empirically, there are switchers in the sample.

multiple $\Delta\tau_{\bar{y}}$ at different \bar{y} as instruments. We suggest constructing a single synthetic average net-of-tax change instrument that is a weighted average of $\Delta\tau_{\bar{y}}$ across the entire income distribution:

$$\sum p\Delta\tau = \sum_{y^j=0}^{\infty} p\Delta\tau(\Delta\tau_{it}) = \sum_{y^j=0}^{\infty} p(y^j)[\tau_{i,t+Dt}(y^j) - \tau_{it}(y^j)]. \quad (28)$$

Empirically, we approximate the income distribution by 200 intervals of USD 1,000 each plus a residual top income interval. We suggest using the weight $p(y^j)$, which is the empirically observed *unconditional* probability density function at y^j . Other weighting functions can be chosen. Using, e.g., $p = 1(y^j = \bar{y})$ results in $\Delta\tau_{\bar{y}}$.

In Appendix A, we show that $\sum p\Delta\tau$ yields a global ETI representing the average elasticity on the taxed, β^{ATT} , in a way that most closely resembles the average treatment effect on the treated (ATT) in the treatment effects literature. $\sum p\Delta\tau$ also has an intuitive interpretation: For a given tax schedule and income distribution, changing the tax schedule so that $\sum p\Delta\tau$ changes by 1% also mechanically (before any behavioral response) changes $\Delta\tau$ by 1%. In the example in Figure 3, $\sum p\Delta\tau = 0.5(\Delta_k\tau^1 + \Delta_k\tau^2)$. Like other instruments using variation within income levels and years, $\sum p\Delta\tau$ groups observations by tax-schedule changes.

In the labor supply literature, the participation net wage between zero and a *constant* hours of work used by Eissa and Hoynes (2004) provides another example of a weighted-average net-of-tax variable that can be used as an instrument (in levels or in changes). Their net-of-tax variable corresponds to $\sum[p = 1(y^j < wage * \bar{h})]t * wage$, where *wage* is the gross wage rate.²⁹

3.5 Policy elasticities

An important limitation of any ETI is that it is measured with respect to the *observed* net-of-tax rate (τ^*). On the other hand, a tax schedule change is expressed in terms of tax rates at different income levels ($\tau(y)$). With progressive tax rates, lowering tax rates at each income level by 1% would decrease observed tax rates by less than 1% as some individuals increasing their income switch to new tax brackets with higher tax rates. The ETI does not incorporate the information on how much different tax rates in different brackets need to be changed for observed net-of-tax rates to change by 1%.

It is possible to define an elasticity of taxable income with respect to each of the mechanical tax-rate variables under policy control or with respect to a function of these tax rates. We call such elasticities policy elasticities (Hendren, 2016), as they directly use the information on how much income changes as one or several tax rates change. The reduced-form estimates of valid instruments yield consistent estimates of such elasticities. Because the

²⁹ While we refer to $\sum p\tau$ in Eq. (28) with $p = p(y^j)$ as “the average net-of-tax rate”, the literature sometimes uses this term as a synonym for the participation net-of-tax up to the observed income level, i.e., for $p = 1(y^j < y^*)$. Because the observed income is endogenous to elasticity heterogeneity, the participation net-of-tax change is not a valid instrument.

first-stage effect of a valid net-of-tax change instrument on observed net-of-tax change is less than one with progressive tax rates, the underlying policy elasticity is greater than the ETI.³⁰

Because there are many tax rates, there are many ways to construct policy elasticities, e.g., the policy elasticity with respect to the base-year net-of tax rate, or the policy elasticity with respect to the first-dollar net-of-tax rate. $\sum p\Delta\tau$ accounts for the policy effects of each of the marginal tax rates. It yields a reduced-form estimate that measures the behavioral taxable income responses relative to the mechanical observed net-of-tax responses in the data. As discussed by Saez et al. (2012), the ratio between behavioral and mechanical *tax revenue* effects is directly related to the marginal deadweight loss, which is an efficiency measure comparable across tax reforms. With linear budget sets without individuals moving between tax brackets, the standard ETI provides a sufficient statistic for efficiency analysis. With nonlinear budget sets, Blomquist and Simula (2016) showed that the marginal deadweight loss is lower when individuals switch tax brackets. Our policy elasticity accounts for such behavioral effects and provides a better efficiency measure for tax reforms generating the type of tax-rate changes seen in the data.³¹

In the labor supply literature, Eissa and Hoynes (2004) estimated the reduced-form effect of the participation net wage on employment rather than using it as an instrument. This yields a policy elasticity with respect to the participation net wage.³²

4. Data and graphical analysis

4.1 Data

We use data from the NBER panel of tax returns from 1979 to 1990, also known as the Continuous Work History File, which is the same data as used in a series of recent papers, e.g., Gruber and Saez (2002), Kopczuk (2005), and Weber (2014). It contains detailed administrative information on taxes and income variables and includes a limited set of demographic background variables. See Gruber and Saez (2002) for a detailed description of the data set.

An important source of variation in the data comes from the Tax Reform Act of 1986 (TRA86), which was the major reform that simplified the tax structure by reducing the number of tax brackets from 15 in 1986 to five in 1987 and just two in 1988, with the top marginal tax rate declining from 50% to 28%. The act also eliminated the second earner deduction and income averaging, and it increased the personal exemption from USD 2,160 in

³⁰ Another way to say this is that switchers attenuate the first-stage effect of the policy tax-rate change; therefore, the structural elasticity contains an attenuation bias unlike the reduced-form elasticity.

³¹ Because taxable income elasticities are only indicative of tax revenue effects with nonlinear budget sets, they cannot generally be a sufficient statistic for efficiency analysis. Liang (2014) developed a method for estimating tax revenue elasticities and showed that, in his application, inferring the marginal deadweight loss from taxable income policy elasticities leads to similar conclusions.

³² Estimates of preference parameters from structural methods are typically used to construct policy elasticities. Blundell and Shephard (2012) provide an example where all net-of-tax rates are increased by 1%, which corresponds to changing a flat consumption tax rate. Furthermore, scaling reduced-form tax-reform estimates, such as the EITC-effect estimate by Eissa and Liebman (1996), by a function of tax rates yields policy elasticities.

1986 to USD 3,800 in 1987 and the standard deduction from USD 3,760 in 1987 to USD 5,000 in 1988.

We construct the individual budget sets by computing marginal tax rates accounting for federal, state, and payroll tax rates at different income levels using NBER-TAXSIM. We vary earnings in steps of USD 1,000, keeping income from other sources fixed. Deductions/itemizations that vary between individuals are accounted for in the construction of the budget sets.

We use two of the most common measures of taxable income previously used in the literature: actual taxable income (almost exactly as technically defined in the tax forms) and broad income. Broad income is a comprehensive definition of gross income that includes, among other things, wage income, interest income, dividends, and business income. Taxable income consists of broad income minus a number of deductions. We use the constant definition of taxable income, as applicable to the year 1990, and include all adjustments that can be computed from the data for all sample years. The taxable income measures are identical to the ones used by Gruber and Saez (2002) and Weber (2014).

One margin of behavioral response to tax reform is to shift income between different sources, e.g., between taxable income and other components of broad income. If income composition is endogenous to tax-rate changes, estimates based on taxable income could be misleading, but those based on broad income would still be largely valid. The taxable income estimates do, however, incorporate the effects of tax avoidance, which is also a margin of interest for policy and welfare evaluations. In the presence of income shifting, Chetty (2009) showed that the marginal deadweight loss depends on a weighted average of the taxable income and broad income elasticities.³³

Our sample selection criteria are similar to the ones in Gruber and Saez (2002). We drop filers that change filing status and observations with missing values. In the baseline specification, we also truncate our data from below by dropping observations with less than USD 10,000 (1990 price level) to avoid issues with filing thresholds.

Following Weber (2014), we use two-year differences for the first-differenced variables.³⁴ When lagged base-year income is used to construct instruments or control functions, we use two-year lags, which allows use of data only with base years from 1981 to 1988.

Table 1 reports means and standard deviations for our main sample. The first section reports the statistics for the dependent variables: changes in taxable income and broad income and for base-year taxable and broad incomes. In the second section of rows, we first report statistics for the main observed net-of-tax change regressor, $\Delta\tau$. We then report the statistics for the instruments: base-year net-of-tax change $\Delta\tau_0$, its predicted version $\Delta\tau_0^{yt}$, its residualized version $\Delta\tau_0^{-yt}$, and for the 2-lag net-of-tax change $\Delta\tau_{-2}$. The instruments are described in Eqs. (20) to (24).

³³ Doerrenberg et al. (2014) showed similar consequences when deductions are endogenous to tax-rate changes.

³⁴ Our results are insensitive to using three-year differences. Some of those results were reported in Kumar and Liang (2015), which is a previous version of this paper.

Table 1. Sample statistics

Variable	Logged variables		Non-logged variables	
	Mean	Std. dev.	Mean	Std. dev.
Δ taxable income	-0.078	0.589	586	34,143
Δ broad income	-0.029	0.297	383	35,747
Base-year taxable income	10.295	0.622	37,325	42,166
Base-year broad income	10.756	0.518	55,262	50,740
$\Delta\tau$	0.029	0.137	1.708	8.061
$\Delta\tau_0$	0.030	0.080	1.678	4.461
$\Delta\tau_0^{yt}$	0.028	0.035		
$\Delta\tau_0^{-yt}$	0.002	0.070		
$\Delta\tau_{-2}$	0.026	0.075	1.480	4.303
$\sum p\Delta\tau$	0.060	0.107	3.432	7.066
$\Delta\tau_{\bar{y}=10}$	0.056	0.137	3.298	9.201
$\Delta\tau_{\bar{y}=25}$	0.061	0.133	3.378	8.633
$\Delta\tau_{\bar{y}=50}$	0.068	0.117	3.743	7.345
$\Delta\tau_{\bar{y}=100}$	0.068	0.121	3.751	7.150
Single-filing singles			0.279	0.448
Joint-filing married			0.668	0.471
Other filing status			0.053	0.224
0 children			0.570	0.495
1 child			0.155	0.362
2 children			0.187	0.390
3 children			0.087	0.282

Notes: Monetary outcomes are in USD at the 1990 price level. Change in tax rate variables are in %-points.

In the third section of rows, we report the statistics for our (density-weighted) average net-of-tax change instrument $\sum p\Delta\tau$ and the net-of-tax changes $\Delta\tau_{\bar{y}}$ at four constant income levels of \bar{y} : USD 10,000, 25,000, 50,000, and 100,000. These instruments are described in Eqs. (28) and (26). While we report logarithmic versions of the variables used in the regressions in columns two and three, we report the non-logarithmic versions in columns four and five to give a better sense of magnitudes, when applicable. In the fourth section of rows, we report the shares of individuals with different filing status and number of children (dependents).

Mean taxable income is USD 37,325, which is considerably lower than mean broad income of USD 55,262. The standard deviations are similar in size to the means. The mean changes in these variables are small, but the standard deviations of the changes are almost as large as their level versions. The mean of $\Delta\tau$ is 1.708 %-points, which is similar to the means of $\Delta\tau_0$ and $\Delta\tau_{-2}$. Standard deviations are four times the mean for the regressor and half of that for the instruments.

The standard deviation of $\Delta\tau_0^{-yt}$ (0.070) is double that of $\Delta\tau_0^{yt}$ (0.035). Contrary to conventional belief, the main variation in tax-rate changes in $\Delta\tau_0$ does not come from variation *across* base-year income levels and years but rather from variation *within* base-year income levels and years.

The means of $\Delta\tau_{\bar{y}}$ are between 3.298 and 3.751 %-points and are increasing in \bar{y} . The standard deviations are all higher than the means, indicating that there is plenty of variation in tax-rate changes at constant income levels. The mean of $\sum p\Delta\tau$ is 3.432 %-points, which is a

weighted average of the different $\Delta\tau_{\bar{y}}$ and closest to the mean change at USD 25,000. In terms of the standard deviation, the $\sum p\Delta\tau$ contains a similar amount of variation as $\Delta\tau_0$.

The demographic variables show that most of the sample consists of married joint filers and filers without children.

4.2 Reduced-form and first-stage relationships

In Figure 5, we explore the structural and reduced-form relationships. We plot local polynomial fits of taxable income change against $\Delta\tau$, $\Delta\tau_0$, Δ_0^{yt} , Δ_0^{-yt} , and $\sum p\Delta\tau$. The x-axis ranges from -0.2 to 0.2 and covers over 90% of the sample for each variable. We use a thick dashed-dotted line for $\Delta\tau$, a thick long-dashed line for $\Delta\tau_0$, and a thick solid line for $\sum p\Delta\tau$. We use thin dashed lines Δ_0^{yt} and Δ_0^{-yt} , with shorter dashes for the latter.




Figure 5. Reduced-form relationships

Figure 5 shows that changes in taxable income and observed net-of-tax rate are negatively correlated (thick dashed-dotted line), reflecting that, with progressive tax rates, random shocks in income changes mechanically lead to observed net-of-tax changes in the opposite direction.

For $\Delta\tau_0$ (thick long-dashed line), the correlation with taxable income change is close to zero. That correlation is negative for Δ_0^{yt} (long-dashed line) and positive for Δ_0^{-yt} (short-dashed line). The negative correlation for Δ_0^{yt} likely indicates mean reversion.

For $\sum p\Delta\tau$ (thick line), the correlation with taxable income change is positive. The slope coefficient in this reduced-form relationship is around 0.5 in the central region with higher

probability density. This coefficient can be interpreted as a policy elasticity estimate (see Subsection 3.5).

In Figure 6, we instead explore the first-stage relationships by moving the observed net-of-tax change to the y-axis. The figure shows that all first-stage slope coefficients between $\Delta\tau$ and the instruments are positive and in the region of 0.5 to 1. The instruments seem to be relevant. The ratio of the reduced-form and first-stage estimates yields a structural ETI estimate in the region of 1 for $\sum p\Delta\tau$.




Figure 6. First-stage relationships

In Figure A1 in Appendix B, we plot the reduced-form relationship for $\sum p\Delta\tau$ in each of the sample years. We also provide the probability density of $\sum p\Delta\tau$. The figure shows that the slope coefficients are positive over most parts of the support in all years.

4.3 Relationships by income levels and years

In Figure 7, we explore the base-year income-by-year variation by plotting the relationship between changes in taxable income and observed net-of-tax rate against base-year income. We do so for the post-reform first-difference generated by TRA86 in 1986 where tax-rate change was increasing in income and for the pre-reform first-difference in 1984 without such variation. We use thick lines for the reduced-form relationships and dashed lines for 1984. The first-stage relationships have the y-axis on the right. We have cut the x-axis at the 5th and 95th base-year income percentiles.




Figure 7. Base-year income-by-year variation in tax-rate changes

The thin lines show a positive slope coefficient for the post-reform change, unlike for the pre-reform change. The difference between slope coefficients yields a first-stage estimate of approximately 0.03 $((0.06-0)/2)$. On the other hand, the thick lines show a positive slope coefficient for the post-reform change, unlike for the pre-reform change, in which the slope coefficient is negative, likely due to mean reversion. The difference between the slope coefficients yields a reduced-form estimate of approximately 0.05 $((0.02-(-0.08))/2)$. The ratio of the reduced-form and first-stage estimates yields a structural ETI estimate of 1.7. Part of the variation in taxable income change across base-year income levels could reflect that high-income individuals respond more to net-of-tax increases received by everybody because they have higher elasticities. The estimate of 1.7 could therefore contain a positive elasticity heterogeneity bias.

If we use a single first-difference approach with only 1986, the reduced-form estimate is 0.01 $(0.02/2)$. This gives an ETI estimate of 0.3 $(0.01/0.03)$. This continuous ungrouped version of the grouping estimate in Feldstein (1995) is much lower than ETI estimates between 1 and 3 found in that paper (with three-year differences). This supports the view that inclusion of minor individual-level variation in tax-rate changes leads to lower estimates. However, accounting for trend heterogeneity with pooled first-differences increases the estimate from 0.3 to 1.7 – similar in magnitude to grouping estimates.

The net-of-tax change instruments used in the literature, believed to mainly exploit income-by-year variation in tax-rate changes, yield ETI estimates ranging from 0.2 to 1.0 (e.g., Auten and Carroll, 1999; Mofitt and Wilhelm, 2000, Gruber and Saez, 2002; Kopczuk, 2005; Weber, 2014), which are much lower than Feldstein's estimates. Saez et al. (2014) mentioned the introduction of individual-level variation and the accounting of trend

heterogeneity as potential explanations for the lower ungrouped estimates. Our estimate of 1.7 shows that, jointly, these explanations cannot reconcile the discrepancy.

However, unlike Feldstein’s method, the ungrouped net-of-tax change instruments also use tax-rate variation within income levels and years (see Eqs. (21) to (23)). In Figure 8, we group observations in 1986 by the sign of $\Delta\tau_0^{-yt}$ and reproduce the reduced-form and first-stage relationships for each group separately. We first reproduce the relationships for the two groups jointly using solid lines. We then use short-dashed lines for the group with low $\Delta\tau_0^{-yt}$ (-) and long-dashed lines for the group with high $\Delta\tau_0^{-yt}$ (+).




Figure 8. Variation in tax-rate changes within base-year income levels and years

At each income level in 1986, we can compute a reduced-form and a first-stage estimate based on the difference between the two groups exploiting only the base-year income-by-year residualized variation in tax-rate changes. At log base-year income 10.5, we see that the first-stage difference between the thin long-dashed and short-dashed lines is approximately 0.12 (0.15-0.03). At the same time, the reduced-form difference between the thick long-dashed and dotted lines is 0.04 (-0.08-(-0.12)). The ratio of the reduced-form and first-stage estimates yields a structural ETI estimate of 0.3. Accounting for the positive elasticity heterogeneity bias (as $\Delta\tau_0^{-yt}$ is more likely exogenous) decreases the ETI estimate (from 1.7 to 0.3).

$\Delta\tau_0$ uses both the variation across and within income levels and years. It should therefore yield an ETI estimate between 0.3 and 1.7. We can compare the scope of variation from the two sources in Figure 8. At log base-year income 10.5, the first-stage difference between the groups with high and low $\Delta\tau_0^{-yt}$ is 0.12. This difference is almost double the first-stage difference of 0.06 (0.11-0.05) between individuals with base-year income of 9.5

and 11.5, respectively. Furthermore, Figures 7 and 8 show that the reduced-form and first-stage relationships for $\Delta\tau_0$ more closely follow those for $\Delta\tau_0^{-yt}$ than those for $\Delta\tau_0^{yt}$. We therefore expect $\Delta\tau_0$ to yield an ETI estimate closer to 0.3 than 1.7.

In Figure A2 in Appendix B, we explore the relationships between $\Delta\tau_0$ and base-year income in 1986 and 1984. We also plot the relationships for $\sum p\Delta\tau$. The positive relationship between base-year income in 1986 and $\Delta\tau_0$ reflects income-by-year variation. On the other hand, $\sum p\Delta\tau$ is uncorrelated with income-year interactions, indicating its plausible validity.

In Figure A3, we show the relationships between $\Delta\tau_0$ and base-year income for each year in the data. The correlation with base-year income is highest in 1986 and 1987.

In Figures A4 and A5, we group $\sum p\Delta\tau$ by state as well as by state-filing-children interactions for each year separately and plot the group-level reduced-form relationships. To account for macroeconomic shocks, we year-demean all observations. The group-level correlations seem to be of the same order as the correlation in Figure 5 that includes the variation in tax-rate changes within groups.

5. Estimation results

5.1 Main IV estimates

In Table 2, we report structural IV estimates of the ETI. In the first section, we report estimates for the following instruments: $\Delta\tau_0$, $\Delta\tau_0^{yt}$, \mathbf{z}_0^{yt} , $\Delta\tau_0^{-yt}$, and $\Delta\tau_{-2}$. In the second section, we report estimates using $\sum p\Delta\tau$ and $\Delta\tau_{\bar{y}}$ at \bar{y} equals USD 10,000, 25,000, 50,000, and 100,000. We also report estimates using all four $\Delta\tau_{\bar{y}}$. The instruments are described in Eqs. (19) to (24), (26), and (28).

In column (1), we provide estimates without covariates. In column (2), we add year-fixed effects, which we keep in columns (3) to (8). In column (3), we add a five-piece spline in base-year income, whereas in column (4), we add a five-piece spline in base-year income lagged two years. In column (5), we include both splines. In column (6), we include additional covariates – dummies for state of residence, filing status, and number of children. In column (7), we replace the splines by year-specific splines. In column (8), we additionally replace the demographic covariates by year-specific covariates.

The raw ETI estimate for $\Delta\tau_0$ in column (1) is negative. It turns positive when accounting for time effects in column (2). It increases to a statistically significant estimate of 0.277 in column (3), which includes the base-year spline to account for trend heterogeneity bias, and it is then quite stable as additional demographic covariates are included in columns (5) and (6). The estimates in columns (5) and (6), of approximately 0.26, are of the same magnitude as the estimate of 0.4 from the preferred specification in Gruber and Saez (2002) using the same instrument, although three-year differences were used in that case. We discuss the estimates in columns (7) and (8) shortly.

Table 2. Structural IV estimates of the ETI

Instrument	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta\tau_0$	-0.064 (0.065)	0.096 (0.072)	0.277** (0.078)	0.140 (0.075)	0.263** (0.077)	0.262** (0.076)	0.183* (0.079)	0.129 (0.082)
$\Delta\tau_0^{yt}$	-0.753** (0.120)	-0.655** (0.205)	1.245** (0.284)	-0.603* (0.268)	1.247** (0.278)	1.285** (0.280)		
z_0^{yt}	-0.811** (0.113)	-0.660** (0.171)	0.987** (0.253)	-0.690** (0.215)	0.970** (0.248)	0.962** (0.247)		
$\Delta\tau_0^{-yt}$	0.170* (0.076)	0.178* (0.077)	0.230** (0.079)	0.192* (0.078)	0.215** (0.078)	0.213** (0.077)	0.183* (0.079)	0.130 (0.082)
$\Delta\tau_{-2}$	-0.198 (0.135)	0.196 (0.206)	0.874** (0.274)	0.385 (0.248)	0.527* (0.258)	0.467 (0.259)	0.115 (0.307)	-0.063 (0.365)
$\sum p\Delta\tau$	0.622** (0.084)	0.834** (0.099)	0.832** (0.098)	0.842** (0.100)	0.819** (0.097)	0.696** (0.093)	0.678** (0.093)	0.691** (0.095)
$\Delta\tau_{\bar{y}=10}$	0.354** (0.096)	0.398** (0.101)	0.393** (0.100)	0.401** (0.101)	0.380** (0.098)	0.276** (0.096)	0.220* (0.097)	0.223* (0.099)
$\Delta\tau_{\bar{y}=25}$	0.765** (0.098)	1.049** (0.117)	1.071** (0.117)	1.067** (0.118)	1.060** (0.116)	0.919** (0.111)	0.904** (0.110)	0.922** (0.113)
$\Delta\tau_{\bar{y}=50}$	1.009** (0.106)	1.356** (0.135)	1.313** (0.130)	1.353** (0.134)	1.300** (0.128)	1.182** (0.124)	1.189** (0.123)	1.205** (0.123)
$\Delta\tau_{\bar{y}=100}$	0.697** (0.137)	1.425** (0.202)	1.324** (0.189)	1.387** (0.197)	1.302** (0.185)	1.330** (0.186)	1.361** (0.178)	1.346** (0.177)
$\Delta\tau_{\bar{y}=10,25,50,100}$	0.643** (0.087)	0.851** (0.103)	0.845** (0.102)	0.857** (0.103)	0.832** (0.100)	0.727** (0.097)	0.723** (0.097)	0.741** (0.099)
Year	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Base-year spline	No	No	Yes	No	Yes	Yes	Yes	Yes
2-lag spline	No	No	No	Yes	Yes	Yes	No	No
Demographics	No	No	No	No	No	Yes	No	No
Year*base-year spline	No	No	No	No	No	No	Yes	Yes
Year*2-lag spline	No	No	No	No	No	No	Yes	Yes
Year*demographics	No	No	No	No	No	No	No	Yes

Notes: Each cell reports an estimate from one regression. Two-year differences are used. The splines each contain five pieces. Demographic covariates include fixed effects for state of residence, filing status, and number of children. Each regression contains 29,085 observations. Standard errors are clustered at the individual level. * $p < 0.05$; ** $p < 0.01$.

The estimates using $\Delta\tau_0^{yt}$ are approximately 1.3 in columns (5) and (6) when a base-year spline is included. Exploiting income-by-year variation in tax-rate changes more flexibly by directly using year-specific base-year income splines as instruments ($\Delta\mathbf{z}_0^{yt}$) yields estimates of approximately 1.0. These estimates are of the same magnitude as the estimate of 1.7 in the graphical analysis in Figure 7 with post-reform first-differences from 1986 and with pre-reform first-differences from 1984. They contain the same type of positive elasticity heterogeneity bias because high-income individuals receiving higher net-of-tax changes also are more elastic.

The estimates using $\Delta\tau_0^{-yt}$ are about 0.20, ranging from 0.170 without covariates in column (1) to an estimate of 0.213 with all non-interacted covariates in column (6). These are of the same magnitude as the estimate of 0.3 in the graphical analysis in Figure 8 using only first-differences from 1986.

Because $\Delta\tau_0^{-yt}$ is uncorrelated with income-year interactions by construction, it is more likely to be independent of preferences and satisfy the identifying assumption. The fact that adding covariates such as the base-year spline and demographic covariates do not affect the estimates is reassuring for instrument validity.

The estimate of 0.26 for $\Delta\tau_0$ is between the estimate of 1.3 for $\Delta\tau_0^{yt}$ and the estimate of 0.21 for $\Delta\tau_0^{-yt}$. While the positive elasticity heterogeneity bias is large in the estimate of 1.3, the estimate of 0.26 is closer to the likely consistent estimate of 0.21.³⁵ This is in line with the graphical analysis in Figure 8 showing that the variation in tax-rate changes is larger within income levels and years than across income levels and years. The analysis shows that the main reason for $\Delta\tau_0$ producing much lower estimates than instruments that exploit only income-by-year variation is because $\Delta\tau_0$ contains substantial income-by-year residualized variation. This pattern can also be seen from estimates on the 1993 tax-rate increase presented in Table 2 of Saez et al. (2012), where $\Delta\tau_0$ yields ETI estimates of 0.143 and 0.237, while a corresponding grouped instrument yields estimates of 0.564 and 0.732. We believe we have convincingly shown that the residualized variation that is uncorrelated with elasticity heterogeneity explains the lower estimates for $\Delta\tau_0$.

The estimates using $\Delta\tau_{-2}$ are considerably higher than the estimates using $\Delta\tau_0$ when including income splines. Standard errors are three times larger, and the ETI estimates are more volatile to the set of covariates used. The estimates with the wider sets of covariates in columns (5) and (6) are approximately 0.5 and lower than the preferred estimate of approximately 0.9 in Weber (2014). However, the pattern whereby replacing base-year income with lagged income in constructing the instrument leads to higher estimates is the same.³⁶ The reason is that the underlying consistent local ETI underweights high-elasticity individuals less as high-elasticity individuals comply more with tax-rate changes at income levels further away from base-year income (see Subsection 3.3). Our estimates are lower than

³⁵ One may argue that precision is too low and that the estimates of 0.26 and 0.21 are not statistically significantly different from each other. We could still assert that there is a positive elasticity heterogeneity bias because the estimate of 1.3 is statistically significantly different from the estimates of 0.26 and 0.21.

³⁶ We could have decomposed $\Delta\tau_{-2}$ into a predicted and a residualized instrument. These instruments yield even more imprecise and volatile estimates, and we omit reporting them here.

those in Weber (2014) because we use a single lagged base-year instrument whereas she used multiple lagged base-year instruments.³⁷

For $\sum p\Delta\tau$, which accounts for and exploits tax-rate changes across the entire income distribution, the ETI estimate is 0.622 without covariates. This estimate is of the same magnitude as the estimate of 1 in the graphical analysis in Figures 7 and 8. It increases to 0.834 when accounting for time effects. The fact that adding income splines to account for trend heterogeneity bias does not change estimates supports instrument validity. The stability of estimates is a result of the underlying tax-rate changes used to construct the instrument, being uncorrelated with income-year interactions, as we saw graphically in Figure A2 in Appendix B. The estimates decrease a bit to 0.696 as additional demographic covariates are included.

In general, precision is much higher when using $\sum p\Delta\tau$ than when using the other instruments. The higher and possibly consistent ETI estimate of around 0.70 compared to the local ETI estimate of 0.21 for $\Delta\tau_0^{-yt}$ reflects that $\sum p\Delta\tau$ appropriately weights all individual responses to tax-rate changes. The large difference suggests that individuals reacting to tax-rate changes at other income levels than the base-year income level have substantially higher elasticities. In comparison, the elasticity heterogeneity bias could only explain the discrepancy between the estimate of 0.26 for $\Delta\tau_0$ and the estimate of 0.21 for $\Delta\tau_0^{-yt}$. Therefore, differences in sets of compliers and the correlation between the degree of compliance and elasticity heterogeneity lead to a substantially larger discrepancy.

$\Delta\tau_{\bar{y}}$ with different \bar{y} yield ETI estimates between 0.276 and 1.330 with income splines and demographic covariates. The estimate increases as \bar{y} increases. The ETI estimate using $\sum p\Delta\tau$ is, in principle, a weighted average of these local ETI estimates. Using several $\Delta\tau_{\bar{y}}$ instruments as an alternative way to exploit multiple tax-rate changes yields an ETI estimate of 0.727 with income splines and demographic covariates, which is close to the estimate of 0.696 for the single synthetic instrument.³⁸

ETI estimates for $\Delta\tau_{\bar{y}}$ with different \bar{y} can be used to approximate β_i at different base-year income levels. It is then possible to simulate the elasticity heterogeneity bias of using variation in $\Delta\tau_0^{yt}$ to investigate whether our theoretical and empirical results are consistent with each other. We run a minimalistic simulation assuming that TRA86 can be approximated by the two-group example with one cross section of first-differences in Figure 2. In Appendix A, we show that τ_0^{yt} yields a simulated ETI estimate of 1.823, which is of the same magnitude as the actual estimate of 1.285 in column (6).

While covariates that are general across years mitigate concerns due to trend heterogeneity, they cannot fully address elasticity heterogeneity that interacts with the year-

³⁷ Weber's (2014, Table 2) estimates show that dropping $\Delta\tau_0$ from the sets of instruments increases the ETI estimates the most, whereas introducing additional lagged instruments has smaller effects. This supports our compliance explanation; see the next footnote.

³⁸ With multiple lagged-income instruments, Weber (2014) obtained ETI estimates between 0.8 and 1.3. She interpreted the higher estimates compared to base-year estimates as the consequence of lagged-income instruments accounting for the trend heterogeneity bias better. We offer a different explanation: Her specifications are close to our multiple-instrument specification and therefore yield similar higher estimates due to a more representative weighting of individuals with different elasticities (compared to base-year instruments). Unlike our instruments, her instruments are, however, not valid, as they do not account for the elasticity heterogeneity bias. Her confidence intervals are also much wider.

specific observed net-of-tax change. Biases due to elasticity heterogeneity can be mitigated by including year-specific covariates (see Subsection 3.2). The estimate using $\Delta\tau_0$ decreases to 0.183 when including year-specific base-year splines in column (7). This estimate is similar to the estimate of 0.213 when using $\Delta\tau_0^{-yt}$ with a general base-year spline in column (6). Both methods remove the income-by-year variation producing the elasticity heterogeneity bias. The estimate decreases a bit when additionally including year-specific demographic covariates.

For $\sum p\Delta\tau$, the elasticity estimate of 0.691 shows that it is remarkably robust to addressing potential remaining elasticity heterogeneity issues by adding year-specific base-year splines and demographic covariates.

5.2 Grouped estimates

Variation in tax-rate changes between demographic groups could possibly be less contaminated by preferences than the variation within demographic groups. Using such group-level variation could therefore serve as a sensitivity test. In comparing grouped and ungrouped estimates, it is, however, important to keep in mind that a discrepancy may reflect both differences in bias and (weighting functions of) the underlying consistent ETIs.

We construct groups based on state of residence, filing status, number of children, and the double and triple interactions between these variables. In Table 3, we report estimates using net-of-tax change instruments averaged within each group. We control for the non-interacted variables. The specification using $\Delta\tau_0$ is similar in spirit to the instrument used by Burns and Ziliak (2016). Columns denote instruments and rows denote groupings.

Table 3. Grouping estimates

Grouping	(1) $\Delta\tau_0$	(2) $\Delta\tau_0^{yt}$	(3) $\Delta\tau_0^{-yt}$	(4) $\Delta\tau_{-2}$	(5) $\sum p\Delta\tau$
State	1.521** (0.402)	1.478* (0.645)	1.537** (0.425)	1.058* (0.472)	0.905** (0.343)
Filing	0.740 (0.479)	-0.044 (0.864)	0.902 (0.551)	0.784 (0.484)	-2.646 (3.642)
Children	1.038** (0.322)	0.787 (0.766)	1.056** (0.325)	1.000** (0.329)	0.938** (0.348)
State-Filing	0.847** (0.263)	1.246* (0.539)	0.760** (0.273)	0.749* (0.363)	0.530* (0.255)
State-Children	0.806** (0.233)	1.351* (0.580)	0.720** (0.239)	0.874** (0.335)	0.825** (0.214)
Filing-Children	0.618 (0.328)	0.118 (0.777)	0.682 (0.349)	0.855* (0.361)	1.224* (0.515)
State-Filing-Children	0.709** (0.209)	1.733** (0.552)	0.571** (0.219)	1.050** (0.341)	0.633** (0.190)

Each cell reports an estimate from one regression. Two-year differences are used. All regressions include five-piece splines in base-year and lagged base-year incomes, as well as fixed effects for state of residence, filing status, and number of children. Each regression contains 29,085 observations. Standard errors are clustered at the individual level. * $p < 0.05$; ** $p < 0.01$.

While ETI estimates using filing groups are imprecise, the other estimates using $\sum p\Delta\tau$ in column (5) vary between 0.530 and 1.224, which are comparable to the preferred ungrouped estimate of around 0.7. Standard errors are large; the estimate of 0.633 using the triple

interaction is the most precise with standard errors that are double those of the ungrouped estimates.

Turning to the base-year and 2-lag instruments in columns (1) to (4), the estimates using $\Delta\tau_0$ are now closer to the estimates using $\Delta\tau_0^{yt}$ and $\Delta\tau_0^{-yt}$. This can be explained by base-year income differences within years being smaller across groups than within groups, which leads to a smaller elasticity heterogeneity bias. The estimates are higher than the ungrouped estimates for $\Delta\tau_0$ and $\Delta\tau_0^{-yt}$. Except for grouping by filing status, they are all between 0.571 and 1.537, and they are not that different from estimates using $\sum p\Delta\tau$. This indicates that the grouped estimates overweigh low-elasticity individuals less than the ungrouped estimates. A potential explanation is that individuals are less likely to move between tax brackets due to the smaller across-group variation in tax-rate changes. More complying high-elasticity individuals staying in the base-year bracket leads to higher estimates.

In Table A1 in Appendix B, we report subsample estimates dropping individuals with more than 60,000 USD and dropping outliers in the dependent variable, the main regressor, and instruments, as well as for each year separately. The main conclusion is that the estimates using $\sum p\Delta\tau$ in column (5) are much more stable across specifications compared to the estimates for the other instruments. Similar to the grouped estimates, these estimates reflect that the correlation between degree of compliance and elasticity heterogeneity vary more across subsamples for the other instruments.

5.3 Reduced-form, first-stage, and broad income estimates

In the first section of Table 4, we report reduced-form and first-stage taxable income estimates, in addition to the structural estimates. The reduced-form estimates yield policy elasticity estimates (see Subsection 3.5). In the second section of Table 4, we report broad income estimates. These are important as the marginal deadweight loss depends on both taxable income and broad income elasticities (Chetty, 2009).

For taxable income, the policy elasticity estimate for $\sum p\Delta\tau$ is 0.455, which is roughly two-thirds of the structural estimate of 0.696. The first-stage estimate is 0.653. To get more clarity, consider changing tax schedules in the same way as they vary in the data. The policy elasticity roughly implies that when such changes lead to a mechanical increase in observed net-of-tax rates of 1% with an associated increase in tax revenues, taxable income increases by 0.455%.

The broad income structural elasticity estimate is 0.082 for $\Delta\tau_0$. The estimate using $\Delta\tau_0^{-yt}$ is similar. The estimate of 0.172 for $\Delta\tau_{-2}$ is larger but not statistically significant. $\sum p\Delta\tau$ yields a broad income elasticity estimate of 0.205. As for taxable income, this instrument yields the highest estimate indicating that the other estimates underweight high-elasticity individuals. $\sum p\Delta\tau$ also yields a broad income policy elasticity of 0.133.

Table 4. Reduced-form, first-stage, and broad income estimates

	(1)	(2)	(3)	(4)	(5)
	$\Delta\tau_0$	$\Delta\tau_0^{yt}$	$\Delta\tau_0^{-yt}$	$\Delta\tau_{-2}$	$\sum p\Delta\tau$
Dependent variable: taxable income change					
IV	0.262** (0.076)	1.285** (0.280)	0.213** (0.077)	0.467 (0.259)	0.696** (0.093)
Reduced form	0.183** (0.052)	1.656** (0.357)	0.148** (0.053)	0.099 (0.053)	0.455** (0.055)
First stage	0.697** (0.015)	1.288** (0.074)	0.695** (0.015)	0.211** (0.017)	0.653** (0.017)
Dependent variable: broad income change					
IV	0.082* (0.041)	0.021 (0.181)	0.086* (0.041)	0.172 (0.137)	0.205** (0.052)
Reduced form	0.058* (0.028)	0.027 (0.233)	0.060* (0.029)	0.036 (0.028)	0.133** (0.032)
First stage	0.701** (0.015)	1.283** (0.074)	0.697** (0.015)	0.210** (0.017)	0.650** (0.017)

Each cell reports an estimate from one regression. Two-year differences are used. All regressions include five-piece splines in base-year and lagged base-year incomes, as well as fixed effects for state of residence, filing status, and number of children. Each regression contains 29,085 observations. Standard errors are clustered at the individual level. * $p < 0.05$; ** $p < 0.01$.

From an efficiency point of view, our taxable income and broad income policy elasticities of 0.455 and 0.133 are the estimates that are the most indicative of the marginal deadweight loss.³⁹

6. Conclusion

We introduced elasticity heterogeneity in the estimation of the ETI in the standard IV setting in first-differences. We showed that elasticity heterogeneity, in addition to trend heterogeneity, is an important source of bias. Instruments used in the literature are invalid because they are by construction contaminated by elasticity heterogeneity. In particular, income-by-year variation in tax-rate changes is endogenous to elasticity heterogeneity. We identified income-by-year residualized and constant-income net-of-tax changes as two types of potentially valid instruments. We also explicitly derived the weights given to different individual elasticities in the different local consistent ETIs that IV could yield. Most instruments do not yield a policy relevant ETI. In particular, we proved that compared with other instruments, the base-year net-of-tax change gives the greatest weight to relatively inelastic individuals.

Based on the insights of our theoretical analysis that identified problems of existing instruments, we proposed constructing an average net-of-tax rate change instrument that weights each constant-income net-of-tax rate change by the income level's empirically observed probability density. This instrument exploits and accounts for effects of changes in

³⁹ These estimates are not that different from the biased ETI estimates of around 0.4 for taxable income and 0.1 for broad income reported by Saez et al. (2002). The reason is that the bias, the overweighting of low-elasticity individuals, and the disregarding of switchers in their estimates lead to effects in different directions that roughly cancel each other out.

the entire tax structure and yields a consistent average ETI on the taxed. Furthermore, this instrument's reduced-form estimate represents a policy elasticity reflective of taxable income and efficiency effects of tax-rate changes in the data.

We illustrated the importance of our theoretical analysis using the NBER tax panel for 1979-1990. The average net-of-tax rate change yielded an ETI of around 0.7. The estimate is robust to inclusion of income splines and demographic controls, and even to inclusion of year-specific versions of these covariates. Furthermore, it is relatively insensitive to the sole use of demographic group-level variation in tax-rate changes for identification. We interpret these results as evidence of instrument validity. We also found a taxable income policy elasticity of 0.46. For broad income, we found an elasticity of 0.21 and a policy elasticity of 0.13.

Reconciling estimates in the literature based on different methods was an important aim of our empirical analysis. We estimated an ETI of 0.26 using the base-year net-of-tax rate change instrument. We then isolated the continuous base-year income-by-year variation and obtained estimates of 1.0 to 1.3. On the other hand, using the base-year income-by-year residualized variation yielded a consistent ETI of around 0.2. The discrepancy between the estimates of 0.2 and 1.0 to 1.3 reflects a large positive elasticity heterogeneity bias. On the other hand, the discrepancy between the consistent base-year ETI of 0.2 and the ETI of 0.7 for our newly suggested average net-of-tax change instrument reflects that the base-year instrument significantly overweights low-elasticity individuals.

In light of our analysis, we have to reinterpret the divergence between estimates in the literature. The grouping estimates (1 to 3 in, e.g., Feldstein, 1995) were larger than the subsequent ungrouped estimates (0.2 to 1.5 in, e.g., Gruber and Saez, 2002; Weber, 2014) because grouping methods exclude tax-rate variation within given income levels and years, and therefore, suffer from a larger elasticity heterogeneity bias. The discrepancies between the ungrouped estimates are, on the other hand, primarily due to differences in how each elasticity is weighted.

Appendices

Appendix A. Derivation of theoretical results

Proof of IV providing a local elasticity β^{LATE} in Eq. (14). Assuming that $\Delta\tau$ is independent of \mathbf{e} , we have the following covariances:

$$\begin{aligned}
\sigma_{E_e(\Delta y|\Delta\tau), E_e(z|\Delta\tau)} &= E_{it}[E_e(\Delta y|\Delta\tau)E_e(z|\Delta\tau)] - E_{it}[E_e(\Delta y|\Delta\tau)]E_{it}[E_e(z|\Delta\tau)] \\
&= E_{\Delta\tau}[E_e(\Delta y|\Delta\tau)E_e(z|\Delta\tau)] - E_{\Delta\tau}[E_e(\Delta y|\Delta\tau)]E_{\Delta\tau}[E_e(z|\Delta\tau)] \\
&= E_{\Delta\tau}\left[E_e(\Delta y|\Delta\tau)[E_e(z|\Delta\tau) - E_{\Delta\tau}[E_e(z|\Delta\tau)]]\right] \\
&= E_{\Delta\tau}\left[E_e(\beta_i\Delta\tau|\Delta\tau)[E_e(z|\Delta\tau) - E_{\Delta\tau}[E_e(z|\Delta\tau)]]\right] \\
&= E_{it}[\beta_i\Delta\tau[E_e(z|\Delta\tau) - E_{it}(z)]],
\end{aligned} \tag{A1}$$

$$\sigma_{E_e(\Delta\tau|\Delta\tau), E_e(z|\Delta\tau)} = E_{it}[\Delta\tau[E_e(z|\Delta\tau) - E_{it}(z)]], \tag{A2}$$

which leads to the equality in Eq. (14). ■

Definition of β^{ATT} . We define it through the weighting function:

$$w_{it}^{ATT} = \frac{\Delta\tau[E_e(\Delta\tau|\Delta\tau) - E_{it}(\Delta\tau)]}{\sum_{it} \Delta\tau[E_e(\Delta\tau|\Delta\tau) - E_{it}(\Delta\tau)]}. \tag{A3}$$

Proof of $\Delta\tau_0$ providing the smallest β^{LATE} in Eq. (27). We can investigate the effect of y^z on β^{LATE} by investigating the effect of y^z on w_{it}^{LATE} . Let us start by investigating the effect of increasing y^z when $y^z \geq y_{it}$. For a data set with a given collection of budget set changes, let us first restrict ourselves to observations with the same y_{it} and τ_{it} in the same year with $\Delta\alpha = 0$. Let $\tau^{y^z} = \tau(y^z)$. For observations with $y_{i,t+Dt} < y_{it}$, $\Delta\tau^{y^z}$ is ceteris paribus uncorrelated with $\Delta\tau$ because $\Delta y < 0$, and changing y^z does therefore not affect w_{it}^{LATE} . We therefore only need to consider observations with $y_{i,t+Dt} \geq y_{it}$ and $\Delta y \geq 0$ for which $\Delta\tau \geq 0$ because $\Delta\alpha = 0$.

Because y_{it} and τ_{it} are fixed:

$$\Delta\tau = \tau_{i,t+Dt}^{y_{i,t+Dt}} - \tau_{it}^{y_{it}} = \Delta\tau(\tau_{i,t+Dt}^{y_{i,t+Dt}}), \tag{A4}$$

$$\Delta\tau^{y^z} = \tau_{i,t+Dt}^{y^z} - \tau_{it}^{y^z} = \Delta\tau^{y^z}(\tau_{i,t+Dt}^{y^z}). \tag{A5}$$

With progressive tax rates, $\Delta\tau^{y^z}$ is stronger correlated with $\Delta\tau$ the closer y^z is to $y_{i,t+Dt}$ because:

$$\frac{d(\tau_{i,t+Dt}^{y_{i,t+Dt}} - \tau_{i,t+Dt}^{y^z})}{dy^z} = \begin{cases} \leq 0 & \text{if } y^z \leq y_{i,t+Dt} \\ \geq 0 & \text{if } y^z > y_{i,t+Dt} \end{cases}, \tag{A6}$$

Now, as y^z changes:

$$\begin{aligned}
\text{sign}\left[\frac{dw_{it}^{LATE}}{dy^z}\right] &= \text{sign}\left[\Delta\tau[\Delta\tau^{y^z} - E_{it}(\Delta\tau^{y^z})]\right] \\
&= \begin{cases} \geq 0 & \text{if } y^z \leq y_{i,t+Dt} \\ \leq 0 & \text{if } y^z > y_{i,t+Dt} \end{cases}.
\end{aligned} \tag{A7}$$

When y^z increases, observations with $y_{i,t+Dt} > y^z$ gain weight at the expense of observations with $y_{i,t+Dt} \leq y^z$. With fixed y_{it} and $\Delta\alpha = 0$:

$$\begin{aligned} y_{i,t+dt} &= y_{it} + \beta_i \Delta\tau \\ \frac{d\beta_i}{dy_{i,t+dt}} &= \frac{1}{\Delta\tau} > 0. \end{aligned} \quad (\text{A8})$$

Observations with $y_{i,t+Dt} > y^z$ that gain weight are therefore relatively more elastic. Therefore, for $y^z \geq y_{it}$, increasing y^z increases β^{LATE} .

For $y^z < y_{it}$, the reverse pattern emerges. Increasing y^z decreases β^{LATE} as relatively less elastic individuals gain weight. It follows that β^{LATE} is minimized when $y^z = y_{it}$. The argument applies for each y_{it} and τ_{it} in each year, and allowing for $\Delta\alpha \neq 0$ is unproblematic as it just adds noise, which proves Eq. (27). ■

$\sum p\Delta\tau$ in Eq. (28) providing the best approximation of β^{ATT} . Let $\tau^{y^j} = \tau(y^j)$. We have:

$$E_e(\Delta\tau|\Delta\tau) = \sum_{y^j=0}^{\infty} \Delta \left[\tau^{y^j} p(y^j|\tau^{y^j}) \right], \quad (\text{A9})$$

$$E_e\left(\sum p\Delta\tau \mid \Delta\tau\right) = \sum_{y^j=0}^{\infty} \Delta \left[\tau^{y^j} p(y^j) \right]. \quad (\text{A10})$$

Among unconditional observed probability densities at different income levels, $p(y^j)$, is clearly that best approximation of the conditional probability density $p(y^j|\tau^j)$. Replacing $E_e(\Delta\tau|\Delta\tau)$ by $E_e(\sum p\Delta\tau \mid \Delta\tau)$ therefore leads to a weighting function that closest resembles the w_{it}^{ATT} is Eq. (A3). Of course, it may be possible to develop a method that estimates $p(y^j|\tau^j)$ and recover β^{ATT} more precisely. Such a method would likely be much more complicated to implement and need distributional assumptions. The setting needed would unlikely nest the just-identified standard first-difference IV setting.

Simulation of the elasticity heterogeneity bias using the example in Figure 2: According to the discussion in connection to Figure 2: $\beta_0^{yt} = (\beta_2\Delta\tau^2 - \beta_1\Delta\tau^1)/(\Delta\tau^2 - \Delta\tau^1) \geq \beta_2 \geq \beta_1$. We simulate β_0^{yt} using the subsample with 1986 as base-year. Dividing it into two groups with the same number of observations by base-year income, we find that the median income is USD 43,000 for the high-income treated group (group 2), and it is USD 18,000 for the low-income control group (group 1). We approximate $\beta_2 = 1.233$ and $\beta_1 = 0.621$ with the ETI estimates obtained using $\Delta\tau_{\bar{y}=43}$ and $\Delta\tau_{\bar{y}=18}$ as instruments, respectively. Furthermore, sample statistics show that $\Delta\tau^2 = 0.116$ and $\Delta\tau^1 = 0.057$. Therefore, $\beta_0^{yt} = (\beta_2\Delta\tau^2 - \beta_1\Delta\tau^1)/(\Delta\tau^2 - \Delta\tau^1) = 1.823$ in the simulation, which is of the same magnitude as the estimated $\beta_0^{yt} = 1.285$ (estimated using the full sample).

Appendix B. Additional empirical results




Figure A1. Yearly reduced-form relationships for the average net-of-tax change




Figure A2. Relationships between instruments and base-year income




Figure A3. The base-year net-of-tax change by base-year income levels and years




Figure A4. Reduced form for the average net-of-tax change by state




Figure A5. Reduced form for the average net-of-tax change by state, filing, and children

Table A1. Subsample estimates

Subsample	(1) $\Delta\tau_0$	(2) $\Delta\tau_0^{yt}$	(3) $\Delta\tau_0^{-yt}$	(4) $\Delta\tau_{-2}$	(5) $\sum p\Delta\tau$
Income < 60,000	0.196* (0.085)	1.449** (0.444)	0.168 (0.086)	0.246 (0.309)	0.675** (0.104)
No outliers	0.330** (0.079)	1.365** (0.293)	0.278** (0.080)	0.548* (0.248)	0.759** (0.102)
1981	-0.020 (0.119)	0.785 (0.850)	-0.028 (0.120)	1.023 (1.091)	0.076 (0.151)
1982	-0.168 (0.221)	-2.034 (1.859)	-0.152 (0.223)	-0.943 (1.050)	0.972** (0.332)
1983	0.405 (0.521)	0.068 (3.892)	0.407 (0.524)	-3.445 (3.815)	0.771* (0.300)
1984	0.837 (0.563)	-1.749 (3.355)	0.835 (0.563)	4.206 (9.236)	0.612* (0.282)
1985	0.410 (0.227)	35.256 (183.814)	0.421 (0.228)	-0.017 (0.979)	0.956** (0.190)
1986	0.115 (0.196)	0.328 (1.457)	0.116 (0.196)	-0.450 (0.487)	0.862** (0.173)
1987	0.190 (0.251)	-0.122 (6.123)	0.191 (0.251)	0.112 (1.343)	1.107** (0.334)
1988	0.109 (0.253)	6.274 (22.127)	0.112 (0.253)	-0.041 (1.146)	1.476** (0.435)

Each cell reports an estimate from one regression. Two-year differences are used. All regressions include five-piece splines in base-year and lagged base-year incomes, as well as fixed effects for state of residence, filing status, and number of children. Standard errors are clustered at the individual level. * $p < 0.05$; ** $p < 0.01$.

References

- Angrist, J., Imbens, G., 1995. Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American Statistical Association* 90, 431-442.
- Auten, G., Carroll, R., 1999. The effect of income taxes on household income. *Review of Economics and Statistics* 81, 681-693.
- Blomquist, S., Kumar, A., Liang, C.-Y., Newey, W., 2014. Individual heterogeneity, nonlinear budget sets, and taxable income. Uppsala Center for Fiscal Studies Working Paper 2014:1, Uppsala University.
- Blomquist, S., Simula, L., 2016. Marginal deadweight loss when income tax is nonlinear. Working Paper, Université de Grenoble Alpes INRA-GAEL.
- Blomquist, S., Selin, H., 2010. Hourly wage rate and taxable labor income responsiveness to changes in marginal tax rates. *Journal of Public Economics* 94, 878-889.
- Blundell, R., Duncan, A., Meghir, C., 1998. Estimating labor supply responses using tax reforms. *Econometrica* 66, 827-861.
- Blundell, R., Shephard, A., 2012. Employment, hours of work and the optimal taxation of low income families. *Review of Economic Studies* 72, 481-510.
- Burns, S., Ziliak, J., 2016. Identifying the elasticity of taxable income. *Economic Journal*, forthcoming, doi: 10.1111/eoj.12299.
- Burtless, G., Hausman, J., 1978. The effect of taxation on labor supply: evaluating the Gary negative income tax experiment. *Journal of Political Economy* 86, 1103-1130.
- Card, D., 2001. Estimating the return to schooling: progress on some persistent econometric problems. *Econometrica* 69, 1127-1160.
- Chetty, R., 2009. Is the taxable income elasticity sufficient to calculate deadweight loss? The implications of evasion and avoidance. *American Economic Journal: Economic Policy* 1, 31-52.
- Chetty, R., Friedman, J., Olsen, T., and Pistaferri, L., 2011. Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: evidence from Danish tax records. *Quarterly Journal of Economics*, 126, 749-804.
- Dagsvik, J., 1994. Discrete and continuous choice, max-stable processes and independence from irrelevant attributes. *Econometrica* 62, 1179-1205.
- Doerrenberg, P., Peichl, A., Siegloch, S., 2014. Sufficient statistic or not? The elasticity of taxable income in the presence of deduction possibilities. ZEW Discussion Paper No. 14-078, Mannheim.
- Eissa, N., Hoynes, H., 2004. Taxes and the labor market participation of married couples: the earned income tax credit. *Journal of Public Economics* 88, 1931-1958.
- Eissa, N., Liebman, J., 1996. Labor supply response to the earned income tax credit. *Quarterly Journal of Economics* 61, 605-637.
- Feenberg, D., and Poterba, J., 1993. Income inequality and the incomes of very high-income taxpayers: evidence from tax returns. In *Tax Policy and the Economy* 7, 145-177, Poterba, J. (ed.), MIT Press, Cambridge and London.
- Feldstein, M., 1995. The effect of marginal tax rates on taxable income: a panel study of the 1986 tax reform act. *Journal of Political Economy* 103, 551-572.

- Feldstein, M., 1999. Tax avoidance and the deadweight loss of the income tax. *Review of Economics and Statistics* 81, 674-80.
- Gruber, J., Saez, E., 2002. The elasticity of taxable income: evidence and implications. *Journal of Public Economics* 84, 1-32.
- Hastie, T., and Tibshirani, R., 1993. Varying-coefficient models. *Journal of the Royal Statistical Society B* 55, 757-796.
- Hausman, J., 1985. The econometrics of nonlinear budget sets. *Econometrica* 53, 1255-1282.
- Heckman, J., Vytlacil, E., 1998. Instrumental variables methods for the correlated random coefficient model: estimating the average rate of return to schooling when the return is correlated with schooling. *Journal of Human Resources* 13, 974-987.
- Hendren, N., 2016. The policy elasticity. *Tax Policy and the Economy* 30, 51-89.
- Hoynes, H., 1996. Welfare transfers in two-parent families: labor supply and welfare participation under AFDC-UP. *Econometrica* 64, 295-332.
- Imbens, G., Angrist, J., 1994. Identification and estimation of local average treatment effects. *Econometrica* 62, 467-476.
- Kawano, L., Weber, C., Whitten, A., 2016. Estimating the elasticity of broad income for high-income taxpayers. Available at SSRN: <http://dx.doi.org/10.2139/ssrn.2852048>.
- Keane, M., Mofitt, R., 1998. A structural model of multiple welfare program participation and labor supply. *International Economic Review* 39, 553-589.
- Kopczuk, W., 2005. Tax bases, tax rates and the elasticity of reported income. *Journal of Public Economics* 89, 2093-2119.
- Keane, M., 2011. Labor supply and taxes: a survey. *Journal of Economic Literature*, 49, 961-1075.
- Kumar, A., Liang, C.-Y., 2015. The taxable income elasticity: a structural differencing approach, Uppsala Center for Fiscal Studies Working Paper 2015:1, Uppsala University.
- Liang, C.-Y., 2014. Distribution-free structural estimation with nonlinear budget sets, Uppsala Center for Fiscal Studies Working Paper 2014:4, Uppsala University.
- Lindsey, L., 1987. Individual taxpayer response to tax cuts: 1982-1984, with implications for the revenue maximizing tax rate. *Journal of Public Economics* 33, 173-206.
- Mirrlees, J., 1971. An exploration in the theory of optimal income taxation, *Review of Economic Studies* 38, 175-208.
- Moffitt, R., 1990. The Econometrics of kinked budget constraints. *Journal of Economic Perspectives* 4, 119-139.
- Moffitt, R., Wilhelm, M., 2000. Taxation and the labor supply decisions of the affluent. In *Does Atlas Shrug? The Economic Consequences of Taxing the Rich* 193-234, Slemrod, J. (ed.), Russell Sage Foundation, New York.
- Navratil, J. 1995. Essays on the Impact of Marginal Tax Rate Reductions on the Reporting of Taxable Income on Individual Tax Returns, Doctoral dissertation, Harvard University.
- Saez, E., 2001. Using elasticities to derive optimal income tax rates. *Review of Economic Studies* 68, 205-229.
- Saez, E., 2003. The effect of marginal tax rates on income: a panel study of “bracket creep”. *Journal of Public Economics* 87, 1231-1258.

- Saez, E., 2004. Reported incomes and marginal tax rates, 1960-2000: evidence and policy implications. In *Tax Policy and the Economy* 18, 117-174, Poterba, J. (ed.), MIT Press, Cambridge and London.
- Saez, E., 2010. Do taxpayers bunch at kink points?. *American Economic Journal: Economic Policy* 2, 180-212.
- Saez, E., Slemrod, J., Giertz, S., 2012. The elasticity of taxable income with respect to marginal tax rates: a critical review. *Journal of Economic Literature* 50, 3-50.
- Slemrod, J., 1996. High-income families and the tax changes of the 1980s: the anatomy of behavioral response. In *Empirical Foundations of Household Taxation*, 169-89, Feldstein, M., Poterba, J. (eds.), University of Chicago Press, Chicago and London.
- Slemrod, J., Kopczuk, W., 2002. The optimal elasticity of taxable income. *Journal of Public Economics* 84, 91-112.
- Triest, R., 1998. Econometric issues in estimating the behavioral response to taxation: a non-technical introduction. *National Tax Journal* 51, 761-772.
- Wald, A., 1947. A note on regression analysis. *Annals of Mathematical Statistics* 18, 586-589.
- Van Soest, A., 1995. Structural models of family labor supply: a discrete choice approach. *Journal of Human Resources* 30, 63-88.
- Weber, C., 2014. Toward obtaining a consistent estimate of the elasticity of taxable income using difference-in-differences. *Journal of Public Economics* 117, 90-103.
- Wooldridge, J., 1997. On two stage least squares estimation of the average treatment effect in a random coefficient model. *Economic Letters* 56, 129-133.