CIRANO
Knowledge into action

IS THE ELASTICITY OF TAXABLE INCOME MOSTLY AN INCOME EFFECT ?

XAVIER DUFOUR-SIMARD
PIERRE-CARL MICHAUD
MICHAEL SMART

CS
2025s-05
WORKING PAPER

CIRANO
Knowledge into action

The purpose of the **Working Papers** is to disseminate the results of research conducted by CIRANO research members in order to solicit exchanges and comments. These reports are written in the style of scientific publications. The ideas and opinions expressed in these documents are solely those of the authors.

Les cahiers de la série scientifique visent à rendre accessibles les résultats des recherches effectuées par des chercheurs membres du CIRANO afin de susciter échanges et commentaires. Ces cahiers sont rédigés dans le style des publications scientifiques et n'engagent que leurs auteurs.

CIRANO is a private non-profit organization incorporated under the Quebec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the government of Quebec, and grants and research mandates obtained by its research teams.

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du gouvernement du Québec, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO Partners – *Les partenaires du CIRANO*

Corporate Partners – Partenaires Corporatifs

Autorité des marchés financiers
Banque de développement du Canada
Banque du Canada
Banque Nationale du Canada
Bell Canada
BMO Groupe financier
Caisse de dépôt et placement du Québec
Énergir
Hydro-Québec
Intact Corporation Financière
Investissements PSP
Manuvie
Mouvement Desjardins
Power Corporation du Canada
VIA Rail Canada

Governmental partners - Partenaires gouvernementaux

Ministère des Finances du Québec
Ministère de l'Économie, de l'Innovation et de l'Énergie
Innovation, Sciences et Développement Économique Canada
Ville de Montréal

University Partners – Partenaires universitaires

École de technologie supérieure
École nationale d'administration publique de Montréal
HEC Montreal
Institut national de la recherche scientifique
Polytechnique Montréal
Université Concordia
Université de Montréal
Université de Sherbrooke
Université du Québec
Université du Québec à Montréal
Université Laval
Université McGill

CIRANO collaborates with many centers and university research chairs; list available on its website. *Le CIRANO collabore avec de nombreux centres et chaires de recherche universitaires dont on peut consulter la liste sur son site web.*

© April 2025. Xavier Dufour-Simard, Pierre-Carl Michaud, Michael Smart. All rights reserved. *Tous droits réservés.* Short sections may be quoted without explicit permission, if full credit, including © notice, is given to the source. *Reproduction partielle permise avec citation du document source, incluant la notice ©.*

The observations and viewpoints expressed in this publication are the sole responsibility of the authors; they do not represent the positions of CIRANO or its partners. *Les idées et les opinions émises dans cette publication sont sous l'unique responsabilité des auteurs et ne représentent pas les positions du CIRANO ou de ses partenaires.*

ISSN 2292-0838 (online version)

Is the elasticity of taxable income mostly an income effect ?*

Xavier Dufour-Simard†, Pierre-Carl Michaud‡, Michael Smart§

Abstract/Résumé

We use variation in marginal tax rates and in tax bracket thresholds at which they apply in order to identify the substitution and income effects of tax reforms. We use a triple-difference estimator that exploits variation from subnational tax reforms, for which behavioral responses to taxes are identified e ven i n t he p r esence o f unobservable shocks to the income distribution. While high-income taxpayers respond more to tax changes, our results suggest this reflects much more the income or salience effects of tax reforms, rather than inherent heterogeneity in substitution effects. We discuss the implications for optimal redistributive tax policies.

Nous utilisons la variation des taux marginaux d'imposition et des seuils des tranches d'imposition auxquels ils s'appliquent afin d'identifier les effets de substitution et de revenu des réformes fiscales. Nous utilisons un estimateur à triple différence qui exploite la variation des réformes fiscales infranationales, pour lesquelles les réponses comportementales aux impôts sont identifiées même en présence de chocs non observables sur la distribution des revenus. Si les contribuables à haut revenu réagissent davantage aux changements fiscaux, nos résultats suggèrent que cela reflète davantage les effets de revenu ou de visibilité des réformes fiscales que l'hétérogénéité inhérente aux effets de substitution. Nous examinons les implications pour les politiques fiscales redistributives optimales.

Keywords/Mots-clés: income and substitution effects, tax salience, optimal progressivity / effets de revenu et de substitution, visibilité de l'impôt, progressivité optimale

JEL Codes : D31, H21, H24

Pour citer ce document / To quote this document

Dufour-Simard, X., Michaud, P.-C., & Smart, M. (2025). Is the elasticity of taxable income mostly an income effect ? (2025s-05, Working Papers, CIRANO.) https://doi.org/10.54932/XPAE6815

---

* This paper was written while Smart was visiting HEC Montreal and Centre interuniversitaire québécois de statistiques sociales (CIQSS). Access to data is through CIQSS and Statistics Canada's Canadian Research Data Centre Network (CRDCN). The views expressed in this presentation are those of the authors, and not necessarily those of CIQSS, Statistics Canada, or its research funding partners.

† HEC Montreal

‡ HEC Montreal, CIRANO

§ Corresponding author: University of Toronto and CESifo

'The rich are different from us.'
'Yes, they have more money.'

– Attributed to Ernest Hemingway

1 Introduction

The elasticity of taxable income (ETI) is a key parameter for evaluating the welfare effects of income tax reforms and designing optimal policy.[1] The existing empirical literature has found substantial evidence of heterogeneity in ETI, with greater average response to tax changes observed among high-income taxpayers.[2] This finding is of considerable policy relevance. If higher-income taxpayers respond more to taxes than others do, then the revenue potential of top-bracket tax increases is smaller at higher incomes, and their deadweight costs correspondingly larger. The implications for optimal tax progressivity are clear. On the basis of their heterogeneous ETI estimates, for example, Gruber and Saez (2002) concluded that the optimal tax system consisted of a negative income tax for low-income individuals, "with generally flat or even declining marginal rates for middle and high income taxpayers."

The greater response of the rich to taxes may reflect their greater work flexibility, knowledge of the tax system, or access to technologies of tax avoidance. But an alternative explanation is the salience of tax reforms. For reforms in the top tax bracket, high-income individuals by definition face a larger mechanical change in tax liability, and so may have greater incentives to respond. In this view, the source of the heterogeneity of the rich is, as in Hemingway's famous quip, simply that "they have more money." Such a response is an income effect rather than a substitution effect of taxation and, in standard public finance models, it has a correspondingly lower deadweight cost. Therefore, distinguishing income and substitution effects of tax reforms is important for policy analysis. In this paper, we estimate income-related heterogeneity in the ETI and examine these opposing views of its source.

Our paper makes three contributions. First, we offer a flexible method for identifying heterogeneous responses to tax reforms by income groups, employing a triple-difference (DDD) estimator that exploits differences in tax reforms that is orthogonal to initial income, while controlling semi-parametrically for group- and time-specific factors that may bias more traditional estimates. Our reduced-form DDD estimates show considerable heterogeneity in responses to top-bracket tax increases, with high-income taxpayers reducing reported income more. Second, the reforms we study involve changes in both top-bracket marginal tax rates and the thresholds at which the rates apply. As we demonstrate, this allows us to identify income and substitution effects of tax changes under mild assumptions. We find that heterogeneity in reduced-form responses is explained by the greater income effects of the reforms on high-income taxpayers, rather than any inherent heterogeneity

[1] For influential contributions, see Feldstein (1999); Saez (2001); Slemrod and Kopczuk (2002).

[2] See for example Gruber and Saez (2002); Kleven and Schultz (2014); Milligan and Smart (2019); Kumar and Liang (2020); Jakobsen and Søgaard (2022).

in substitution effects. Third, we explore the policy implications of our empirical findings. We show that the finding of a positive income effect on reported income has substantial implications for the marginal deadweight cost of taxation and the directions of revenue- and welfare-increasing reforms to the tax system.

Our research builds on recent innovations in the estimation of ETI. The empirical literature has long employed taxpayers' lagged income as a predictor of who is treated by tax reforms. But lagged income is confounded with mean-reverting shocks to individual incomes, with trend changes in income inequality, and with income-related heterogeneity in the ETI (Kumar and Liang, 2020), all of which lead to inconsistency of the traditional IV estimator. The existing empirical literature assumes that shocks to the income distribution are stable over time, and controls for them with either parametric income splines (e.g. Weber, 2014) or semi-parametric methods (e.g. Jakobsen and Søgaard, 2022).

Our triple-difference estimator relaxes the identification assumptions behind the existing estimates. We examine reforms in top-bracket tax rates in four large Canadian provinces, which occurred in a staggered fashion over the 2013-16 period. We relax the identifying assumption in Jakobsen and Søgaard (2022) and allow for country-wide macro shocks to impact the relationship between the level of income and changes in income. Hence, we compare the change in reported income of those in provinces who experienced a change in marginal tax rates to the same change for taxpayers with similar pre-reform incomes in other provinces. Non-reforming provinces serve as controls for national shocks to the income distribution in reform years, which may bias traditional DD estimates. Non-reform years serve as controls for province-specific patterns of mean reversion.

In so doing, we build on recent research that proposes identifying the ETI using only tax variation that is orthogonal to lagged income (Burns and Ziliak, 2016; Kumar and Liang, 2020). Unlike those papers, however, we estimate both the substitution and income effects of reforms. Our estimation strategy uses variation in the tax bracket threshold at which tax reforms apply, generated by time-varying provincial differences in tax schedules, in order to identify heterogeneous substitution and income effects of tax changes. To develop an intuition for our method, consider two reforms that induce the same change in the top marginal tax rate, one beginning at a low threshold and the other at a high threshold. Both reforms induce the same substitution effect on reported incomes, but the mechanical effect of the reform on tax liability is larger for the reform that begins from the lower threshold. Therefore, tax rate reforms with different thresholds have different income or salience effects on taxpayers, even when they have identical substitution effects.

Estimates from the stacked triple difference specification suggest substantial responses to the top-bracket tax increases, with the magnitude of responses generally increasing with lagged income. Our estimates indicate that much of the response is due to the income effect of the reforms. We estimate that a one percentage point mechanical increase in a treated taxpayer's average tax rate due to reform causes a 2.4 percent in reported income through the income effect channel. Since income effects are larger at higher income levels, this in turn has implications for the heterogeneity in *substitution* effects. If the income effects of the reform constrained to zero, the estimated compensated elasticities of taxable income are 0.10 for incomes below $200,000 (the 98.5th percentile of the distribution), and 0.98 above that quantile. Thus our results are consistent with the previous literature in finding substantial heterogeneity of responses among high-income taxpayers. When the income effects of the reform are estimated from provincial threshold differences, elasticity is reduced considerably. In our preferred specification, the estimated substitution elasticity is 0.17 for lagged incomes below $200,000, and 0.59 above it.

While our estimates, consistent with previous research, display substantial income-related heterogeneity in the uncompensated ETI, we find that heterogeneity is explained by income effects rather than heterogeneous substitution effects of reforms. At first glance, this result is surprising. Our results imply that the income effect on reported income is positive, whereas the workhorse model of labor supply (e.g. Gruber and Saez, 2002) predicts a negative income effect. We develop a model of tax avoidance that can rationalize a positive income effect in cases where pure avoidance responses dominate labor supply responses. Alternatively, we show that our findings can be rationalized by a model of tax salience and inattention, in which taxpayers facing larger changes in average tax rates due to reforms are more inclined to respond to them.

There is little consensus in the empirical literature on the magnitude or even the sign of the income effect on reported taxable income. Some previous research has found insignificant or small negative income effects of tax reforms (Gruber and Saez, 2002; Kleven and Schultz, 2014). Much of the empirical literature on the elasticity of taxable income has not controlled for virtual income changes of tax reforms (Burns and Ziliak, 2016), even when examining the income-related heterogeneity in responses to tax rate changes.[3] However, a recent literature has emphasized that large tax changes may be especially salient to taxpayers and so induce larger behavioral responses. For example, Feldman et al. (2016) find that taxpayers respond to the loss of a lump-sum tax credit by increasing labor income, which is broadly consistent with the income effects we estimate.

The distinction between income and substitution effects of tax reforms is relevant to evaluating the welfare effects of income taxes. If response to a tax change is through income rather than substitution effects, then a local change in taxes would have no marginal deadweight cost. The resulting behavioral change simply reflect the transfer of resources from taxpayer to the treasury and are identical to what would occur through a non-distortionary lump-sum tax. We show how our estimated income effects lead to different conclusions about the desirability of reforms increasing marginal rate progressivity of taxes.

2 Income and salience effects of tax reforms

We begin by laying out a simple extension of the standard model of behavioral responses to tax reforms that incorporates both labor supply and pure tax avoidance responses. In the model, the taxpayer chooses total income z, and the portion s of income that is received

[3] See for example Kumar and Liang (2020) and Jakobsen and Søgaard (2022).

in an untaxed form. Taxable income $y = z - s$ is taxed according to a general non-linear tax function $T(y)$. Income s may be exempt from taxation by design as in Feldstein (1999), or it may represent income that is unreported to the tax authorities, as in Chetty (2009). The taxpayer maximizes an utility function $U(c, y, s)$ with $U_c > 0$ and $U_y, U_s < 0$ subject to the budget constraint $c = y - T(y) + s$. Thus $-U_y$ and $-U_s$ are the marginal costs of earning income in taxable and non-taxable form respectively. The first-order conditions describing optimal tax responses are $-U_y/U_c = 1 - T'(y)$ and $-U_s/U_c = 1$.

To model behavioral responses to taxation in a way that is independent of the non-linearities in the tax function, we linearize the budget constraint in the neighborhood of the optimum (Saez, 2001), defining the taxpayer's marginal tax rate $\tau = T'(y)$ and virtual income $I = (\tau - a)y$, where $a = T(y)/y$ is the average tax rate. (See Figure 1.) The taxpayer's taxable income function $y^*(1-\tau, I)$ then solves

$$\max U(c, y, s) \;\text{ s.t. }\; c = (1-\tau)y + I + s \tag{1}$$

Consider an arbitrary small tax reform which changes the tax function T. The resulting change in virtual income at y is $dI = (d\tau - da)y$ where $da = dT/y$ is the mechanical change in average tax rate induced by the reform for someone reporting taxable income y. Differentiating $y = y^*(1-\tau, I)$ and dividing by y,

$$\frac{dy}{y} = -\epsilon^u \frac{d\tau}{1-\tau} + \frac{\partial y^*}{\partial I}(d\tau - da) \tag{2}$$

where $\epsilon^u = \partial \log y^* / \partial \log(1-\tau)$ is the uncompensated ETI. Substituting the Slutsky decomposition[4] $\epsilon^u = \epsilon^c + \eta$ where ϵ^c is the compensated ETI, and $\eta = (1-\tau)y^*_I$ is the income effect on taxable income into (2) yields the structural model

$$\frac{dy}{y} = -\epsilon^c \frac{d\tau}{1-\tau} - \eta \frac{da}{1-\tau} \tag{3}$$

Equation (3) generalizes the decomposition in Gruber and Saez (2002) to a setting in which taxpayers may adjust pure avoidance activities as well as labor supply in response to taxes. It relates changes in marginal and average tax rates due to tax reforms to observable changes in taxable income, and it is the basis of our estimation strategies, discussed below.

In this model, the sign of the income (or average-tax) effect on taxable income η is unrestricted, so that a lump-sum tax increase might cause taxable income either to fall or to rise. To see this, consider two special cases of the general model:

[4] The dual to (1) is

$$e(1-\tau, u) = \min c - (1-\tau)y - s \;\text{ s.t. }\; U(c, y, s) \geq u$$

with $e_1 = -y^c(1-\tau, u)$. Since $y^c(1-\tau, u) \equiv y^*(1-\tau, e(1-\tau, u))$, the Slutsky decomposition is

$$\frac{\partial y^*}{\partial(1-\tau)} = \frac{\partial y^c}{\partial(1-\tau)} + y\frac{\partial y^*}{\partial I}$$

or equivalently, after multiplying by $(1-\tau)/y$, $\epsilon^u = \epsilon^c + \eta$.

1. In the first case, studied by Gruber and Saez (2002), there is no pure avoidance activity, so that $s = 0$, and utility is additively separable: $U(c, y, s) = u(c) - g(y)$. In this case, behavioral responses to tax reforms are labor supply responses, in the sense that the sole means of increasing consumption is by increasing taxable income. Given additive preferences, leisure is a normal good in this model. Thus $\eta \leq 0$, and a lump-sum increase in taxes $da > 0$ causes taxable income to rise.

2. In the second case, labor supply is fixed, so that $y = \bar{z} - s$, and utility is $U(c, y, s) = u(c) - h(s)$. In this case, behavioral responses to reforms are pure avoidance effort, such as the substitution of exempt income or deductible consumption for taxable income, or non-reporting of income. With separable preferences, it is easily verified that $\eta \geq 0$ in this model, so that a lump-sum increase in taxes causes taxable income to fall.[5]

These are simply two polar cases of the general behavioral responses captured by (3). A general point that emerges is that, when taxpayers respond to tax reforms more by adjusting avoidance behavior than labor supply, it is reasonable to expect the income effect to be positive rather than negative, so that an increase in taxes paid due to a reform causes a *reduction* in reported taxable income.

Heterogeneity and income effects. In this model, income effects on behavior naturally give rise to heterogeneous responses to top-bracket tax reforms that are correlated with potential income, even in the absence of any differences among taxpayers in underlying preferences or access to tax avoidance technologies. To see this, consider a top-bracket tax reform, i.e. a change in the marginal tax rate $\Delta\tau$ that applies to incomes $y \geq K$. The resulting change in virtual income is $\Delta I = 1(y \geq K)\Delta\tau K$. The resulting mechanical change in the average tax rate is $\Delta a = 1(y \geq K)\Delta\tau(1 - K/y)$. For a given marginal tax rate change, when the taxpayer's income is larger, the mechanical effect on the average tax rate is smaller, and the virtual income effect is larger. The result is heterogeneity in responses to the reform, the magnitude of which is either positively correlated with income (when $\eta > 0$) or negatively correlated with income (when $\eta < 0$).

Given a single top-bracket tax reform, the income effect cannot be identified separately from heterogeneous substitution effects that are correlated with income. What is therefore required is variation in the income effect that is uncorrelated with a taxpayer's potential income. Consider two top-bracket tax reforms $i = L, H$ that induce that same change in marginal tax rate $\Delta\tau$ for incomes above thresholds $K_H > K_L$. This is depicted in Figure 1. By construction, the virtual income change is larger for the reform that applies at the higher threshold, $\Delta I_H > \Delta I_L$. Thus, in principle, income effects can be identified by comparing the response of a taxpayer of a given potential income to two reforms in which tax changes apply at different thresholds. In our application, that variation comes from different jurisdictions that implement tax reforms at different thresholds.

[5]For these preferences, the optimal avoidance satisfies $h'(s^*)/u'(c_0 + \tau s^*) = \tau$. Since u is concave, the income effect on avoidance is negative, and the income effect on taxable income is positive.

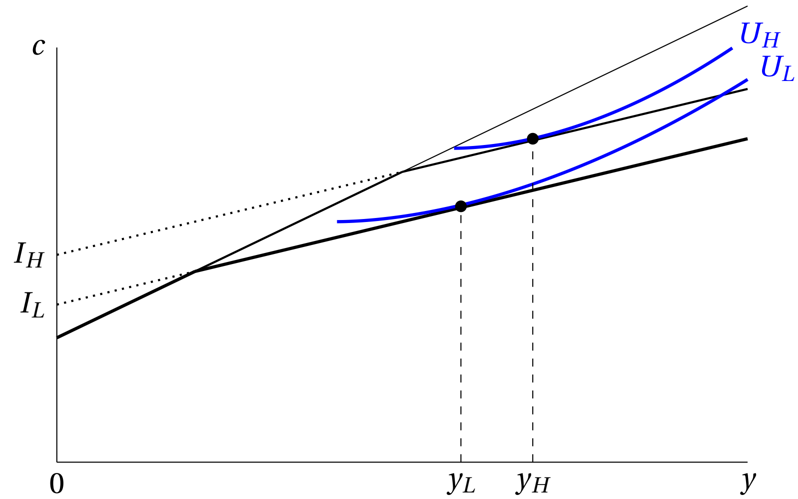

1

Figure 1: Income effects of a tax reform. The figure depicts effects of a given increase in the top marginal tax rate $\Delta\tau$ at two different bracket thresholds $K_H > K_L$. Linearized budget constraints are depicted as dotted lines. For taxpayers with initial income $y > K_H$, the resulting change in virtual income is $\Delta I_i = \Delta\tau K_i$. The reform with the high threshold causes a larger increase in virtual income, $\Delta I_H > \Delta I_L$ and, if the income effect is positive, a smaller reduction in taxable income, $-\Delta y_H < -\Delta y_L$.

Salience vs. income effects. This model suggests that a mechanical rise in a taxpayer's tax liability can cause an increase in tax avoidance behavior. It is important to note that this mechanism is essentially indistinguishable in our approach from a broader model of inattention and salience effects of tax reforms.

To see this, let us suppose that a fraction $p(z)$ of taxpayers with total income z are informed about a tax reform, with the remainder ignorant.[6] An informed taxpayer experiences no income effects, so that the resulting change in income is $\Delta \log y_i = \epsilon^* \Delta \log(1-\tau_i) + u_i$. In contrast, an uninformed taxpayer disregards the tax reform entirely. Thus observed responses to the reform are $E\Delta \log y_i = p(z_i)\epsilon^* \Delta \log(1-\tau_i)$. The probability that a taxpayer is informed and so responds to the reform depends on the salience of the reform, so that

$$p(z) = 1 - \alpha \frac{k}{z} = (1-\alpha) + \alpha \left(1 - \frac{k}{z}\right) = (1-\alpha) + \alpha \frac{\Delta a}{\Delta \tau}$$

where $\alpha \leq 1$ is a parameter. Observed responses to the reform are then a mixture distribu-

[6] A related model of tax salience was considered in Chetty et al. (2013). Other models of imperfect information and non-optimizing behavior would give similar results. For example, under the "ironing" heuristic (Rees-Jones and Taubinsky, 2020; Craig and Slemrod, 2024), taxpayers behave as though their average tax rate was their marginal tax rate. This behavior would result in similar effects of Δa as in the income effects model. For a method to distinguish income and salience effects, see Feldman et al. (2016). [**Perhaps this note should be in the Introduction.**]

tion with conditional expectation

$$E\Delta \log y_i \equiv \epsilon^c \Delta \log(1-\tau_i) - \eta_i\, \Delta a_i$$

where $\epsilon^c = (1-\alpha)\epsilon^*$ is the weighted average ETI, and $\eta_i = \alpha/(1-\tau_i)$ is a parameter measuring the average salience of the reform for treated taxpayers.

Thus we conclude that, with data on a series of top-bracket tax reforms, our structural model of income effects in response to reforms cannot be distinguished from a prominent alternative model of tax salience. In what follows, we use the terms income effect and salience effect interchangeably to describe behavioral responses to tax reforms that are attributable to the mechanical change in tax liabilities that a taxpayer experiences from the reform.

3 Institutions and Data

3.1 Income taxation in Canada

Canada is a federation in which income tax powers are co-occupied by the federal government and the governments of the ten provinces. Provincial taxes are not deductible for federal taxation. Constitutionally, the provinces have wide latitude in designing their personal income tax systems, and they collect substantial revenue from them. In 2022, provincial personal income tax revenues were $132 billion, or 39% of combined federal–provincial revenues.[7]

By agreement, the federal revenue agency collects income taxes on behalf of nine provinces as well as the federal government, with only Quebec operating its own income tax system. Under these tax collection agreements, the agreeing provinces are required to adopt the federal income definition. This common tax base is useful both for estimation and for interpretation of the sensitivity of taxable income to tax rates. However, provinces retain complete flexibility in determining the bracket structure and associated statutory marginal tax rates.[8] Provinces apply progressive rate structures to taxable incomes. The provincial tax rates for high earners vary across the country, ranging from a low of 10 per cent in Alberta before 2016 to a high of 21 per cent in Quebec after 2012. The federal income tax rate for the highest bracket is 33 percent since a 2016 reform.

Provinces have reformed top bracket rates frequently in the past decade. Table 1 reports the marginal tax rate changes and the thresholds at which they applied for reforms in the four largest provinces. In our empirical analysis, we study reforms in the four largest provinces (Ontario, Quebec, Alberta and British Columbia), which increased top marginal tax rates by two to five percentage points in a staggered fashion between 2013 and 2016. The thresholds at which the new rates applied varied between $100,000 and $300,000,

[7]Data from Statistics Canada, as recorded at https://financesofthenation.ca/real-fedprov/.

[8]Provincial government gained these powers in a 2001 reform. For more details, see Milligan and Smart (2019).

Province	First top threshold[a]	MTR change	Second top threshold[a]	MTR change
Quebec (2013)	$100,000	1·8	–	–
Ontario (2014)	$150,000	1·6	$220,000[b]	3·1
British Columbia (2014)[c]	$150,000	2·1	–	–
Alberta (2016)[d]	$125,000	2·0	$150,000	3·0

[a]Thresholds reported in Canadian dollars, equal to approximately US $0.80 at purchasing power parity

[b]Threshold reduced from $500,000, established in 2012

[c]Abolished in 2016-17

[d]Additional MTR increase of 4 points over $200,000 and 5 points over $300,000

Table 1: Top bracket tax reforms. The table shows the marginal tax rate changes in percentage points for the four reforming provinces, and the bracket thresholds of taxable income at which they apply.

affecting between 3.0 and 6.4 per cent of taxpayers, depending on the province. In 2016, a new top federal rate of 33 per cent on incomes over $200,000 was introduced, replacing the previous top federal rate of 29 per cent, bringing the average top rate to 52 percent. As a result of federal and provincial changes, the top take-home rate (one minus the marginal tax rate) has decreased by 13.5 percent on average since 2012. The thresholds at which marginal tax rates rose vary from $100,000 in Quebec to a maximum of $300,000 in Alberta. Two provinces, Ontario and Alberta, adopted reforms with multiple new top tax brackets. Marginal tax rates fell slightly in all provinces around $130,000 due to the indexation of federal bracket thresholds.

3.2 Data sources and definitions

The data for our analysis are from the T1 Personal Master File, a longitudinal database on the universe of taxpayers, created by Statistics Canada from administrative tax records. We exclude taxpayers residing in the northern territories and in four small provinces in the Atlantic region, which experienced both tax increases and tax decreases in the top bracket during our analysis period, yielding a sample of approximately 24.8 million observations per year. We measure changes in incomes and in tax rates over four-year intervals[9], using observations in four years around the data of the tax reform in each reforming province, excluding the year prior to the reform in order to abstract from anticipation effects. We exclude taxpayers who resided in a different province four years earlier, since different provinces have different tax rates. This reduces the sample to approximately 21 million

[9]We winsorize the income change at the first percentile to limit the impact of extreme values on our results.

per year. Since we study reforms at thresholds between $100,000 and $300,000, we restrict the sample to taxpayers with lagged taxable income in the range of $75,000 to $1 million. We also restrict age to be in the range of 18 to 64. The restrictions limit the sample to approximately 1.8 million taxpayers per year (for a year where all provinces are included). Taken together, these sample selection rules result in an analysis sample with approximately 7.6 million taxpayer-year observations. Summary statistics are reported in Table 2 for the analysis subsample. The mean income in the sample is $112,000, which was approximately the 95th percentile of individual incomes in our sample period. Within this high-income sample, individuals are significantly more likely to be male and married than the general population.

Different measures of income might be employed to measure response to taxes. For most of our analysis, we examine changes in Total (or Line 150) Income as recorded on the personal tax form. This is a relatively broad measure of gross (pre-deduction) income that is subject to taxation. As such, it excludes certain income sources that are exempt from taxation, which are chiefly forms of government transfer income such as cash welfare payments and refundable tax credits. Since our analysis sample consists of taxpayers in the upper deciles of the income distribution, these exclusions are unlikely to affect our results.

In some specifications, we also examine results for two narrower measures of income: Taxable Income and T4 Income. Taxable (or Line 260) Income equals Total Income less certain deductions and exemptions. There are relatively few allowable deductions for tax purposes in the Canadian tax system, with most deductions having been replaced in a 1987 reform by non-refundable tax credits (which affect taxes paid but not taxable income). In the 2006-17 period we study, taxpayers declared $14.7 trillion in Total Income, of which $13.1 trillion or about 89 per cent was included in Taxable Income. The share of deductions from Total Income is very similar in our analysis sample (see Table 2). Nevertheless, contributions to registered individual retirement accounts and child care expenses of working parents and some other employment expenses remain deductible, up to some limits. While the ETI for Total Income is likely more policy-relevant than that for Taxable Income (Doerrenberg et al., 2017), we report both. (We use ETI to refer to both the elasticities of total and taxable income. Unless otherwise specified, the ETI estimates reported in this paper are for the broader Total Income measure.) T4 (or Line 101) Income is employment income for which third-party reporting (and typically income tax withholding) is provided to the Canada Revenue Agency. Note that T4 Income excludes income from self-employment, but it includes wages from corporations that may be closely held by the taxpayer or by a family member.

Sample	Mean	S.d.	P_{25}	Median	P_{75}
Total Income (×1000)	142	173	89	112	152
Taxable Income (×1000)	129	167	80	101	136
T4 Income (×1000)	107	126	58	99k	131k
Δ log(Total Income)	-7.4	70.6	-19	5.9	20.6
Δ log(Taxable Income)	-9	71.8	-21.4	4.3	20
Δ log(T4 Income)	-8.4	94.9	-10	7.5	20.6
Male	72%	-	-	-	-
Married	79%	-	-	-	-
Age	49	9	43	50	56

Table 2: Summary statistics. Shows summary statistics for the analysis sample of 7.6 million observations of taxpayers age 18-64 with lagged income in excess of $75,000 and below 1M. Incomes are reported in Canadian dollars, with 1 CAD ≈ 0.74 USD. Income changes are in log points relative to four years prior.

4 Estimation

4.1 Empirical specification

We seek to estimate behavioral responses to changes in taxation from longitudinal data for taxpayers observed in different jurisdictions (provinces). These taxpayers may have experienced changes in marginal tax rates over time. The basic estimating equation following (3) can be written as

$$\Delta \log y_{it} = \alpha + \epsilon^c \Delta \log(1-\tau_{it}) + \gamma \Delta a_{it} + u_{it} \tag{4}$$

where i denotes a taxpayer, t a tax year, and the Δ operator indicates four-year differences, i.e. for a variable x_t, $\Delta x_t = x_t - x_{t-4}$. The term $\log(1-\tau_{it})$ is the log of the take-home rate (one minus the marginal tax rate). The parameter ϵ^c is the compensated elasticity as defined in (3). The term Δa_{it} is the change in the average tax rate and $\eta = -(1-\tau_{it})\gamma$ is the income effect parameter as defined in (3). The error term u_{it} in (4) captures variation in taxable income y_{it} which is not due to behavioral responses to changes in the tax system.

Estimating (4) by least-squares will yield biased estimates of both ϵ^c and γ. Both τ_{it} and a_{it} are calculated using current income y_{it} and therefore are correlated with u_{it} by construction. Let $\tau_{it} = \tau_t(y_{it})$ and $a_{it} = a_t(y_{it})$ reflect that the construction of both these tax variables involve variation over time in the tax function (t) but also depend on the level of reported income y_{it} with non-linear tax schedules. To break the simultaneity, one can use y_{it-4} to compute the (simulated) change in the take-home rate and the average tax rate. The simulated counterparts are $\tilde{\tau}_{it} = \tau_t(y_{it-4})$ and $\tilde{a}_{it} = a_t(y_{it-4})$. If u_{it} is serially uncorrelated, one can use the instruments, $\Delta \log(1-\tilde{\tau}_{it})) = \log(1-\tau_t(y_{it-4})) - \log(1-\tau_{t-4}(y_{it-4}))$ and $\Delta \tilde{a}_{it} = a_t(y_{it-4})) - a_{t-4}(y_{it-4})$ to obtain unbiased estimates of ϵ^c and γ. Estimation is done by two stage least squares.

There are two empirical problems with this strategy. First, $\log y_{it}$ may be serially correlated, absent changes in the tax schedule. In this case, $E \log y_{it-4} u_{it} \neq 0$ and the IV estimator is biased. This happens for example when $\log y_{it}$ follows a persistent mean-reverting process or if there are secular trends in inequality. Past levels of income are informative of future changes in income.

Second, there needs to be independent variation in both the take-home rate and the average tax rate to secure identification of both elasticities. At first glance, heterogeneity in the uncompensated elasticity by lagged income y_{it-4} help identify γ, the income elasticity, under the maintained assumption that ϵ^c is constant. But as emphasized by Kumar and Liang (2020), this may not be the case as heterogeneity may well exist and is likely a function of income levels. To secure identification of both elasticities while allowing for heterogeneity, one needs changes in the tax schedule which generate independent variation in τ_{it} and a_{it} given y_{it-4}.

4.2 Control function approach

To address the first problem, Jakobsen and Søgaard (2022) propose using a control function approach. Absent any change in tax schedules, assume that

$$E(u_{it} | \log y_{it-4}) = f(\log y_{it-4}) + \lambda_t. \tag{5}$$

This assumption states that the relationship between $\log y_{it-4}$ and changes in income absent changes to the tax schedule (u_{it}) is captured by some time-invariant function $f(\cdot)$ and an additive time fixed effect. An income process with a stable mean-reversion process generates a negative relationship between $\log y_{it-4}$ and changes in income (u_{it}). The relationship is time-invariant except for some parallel shift captured by λ_t. This allows for a macro shock that affects the entire distribution equally. The assumption allows to exploit particular changes to the tax schedule. If the tax change affects some identified region of the income distribution, any change in the relationship between lagged income and changes in income can be attributed to a behavioral response to the change in taxes. Jakobsen and Søgaard (2022) approximate non-parametrically the function $f(\cdot)$ by a series of (fixed) bin dummies collected in the vector $\mathbf{D}_{it-4}$. This is an extension of the approach taken in the literature using flexible (e.g. spline or polynomials) controls for $\log y_{it-4}$ (Gruber and Saez, 2002).The estimating equation in (4) becomes

$$\Delta \log y_{it} = \alpha + \epsilon^c \Delta \log(1 - \tau_{it}) + \gamma \Delta a_{it} + \mathbf{D}'_{it-4}\beta + \lambda_t + v_{it} \tag{6}$$

where v_{it} is now an error component which does not depend, by construction, on $\log y_{it-4}$. Under this assumption, the simulated instruments $\Delta \log(1 - \tilde{\tau}_{it})$ and $\Delta \tilde{a}_{it}$ are now immune to the mean-reversion problem provided they contain independent variation conditional on $f(\cdot)$ and the time fixed effect. The maintained identifying assumption is that there is no change in the time invariant relationship $f(\cdot)$ and that time effects are only additive. This precludes for example macro shocks which affect differentially different parts of the income distribution. As suggested by Jakobsen and Søgaard (2022), this suggests a difference-in-difference estimator which lends itself to a visualization of the responses. Suppose there are two four-year differences available ($t = 1, 2$). Suppose there is a change in the marginal tax rate above some threshold in period 2. Denote $POST_t$ as dummy equal to 1 in the year when the tax change occurs (and zero otherwise). Under the identifying assumption (5), we have that

$$\Delta \log y_{it} = \lambda POST_t + \mathbf{D}'_{it-4}\beta + (POST_t \times \mathbf{D}_{it-4})'\theta + v_{it} \tag{7}$$

where θ is a vector of coefficients on the interactions between bin dummies and the POST dummy. If there is a behavioral response to the change in taxes, one would expect the θ estimates in the treated part of the income distribution to be statistically different from zero while they should be close to zero in the untreated part of the income distribution. The treated part of the income distribution can be uncovered by estimating a similar difference-in-difference regression for $\Delta \log(1-\tau_{it})$ (or Δa_{it}) as the dependent variable.

4.3 Exploiting provincial variation

We can exploit policy variation across provinces to (i) relax the identifying assumption in Jakobsen and Søgaard (2022) and (ii) enable the identification of ϵ^c and γ while allowing for heterogeneity in ϵ^c. First, we have provinces who implement changes in top rates at different moments in time. We also have a set of provinces who never change their top rates. This means that we can extend the reduced-form difference-in-difference specification to a third difference (DDD) specification. Let p be an indicator for the province of respondent of taxpayer i. We replace the identifying assumption in (5) by

$$E(u_{it}|\log y_{it-4}) = f_t(\log y_{it-4}) + g_p(\log y_{it-4}) + \lambda_{p,t}. \tag{8}$$

This assumption allows for country-wide changes over time in the mean-reversion relationship $f(\cdot)$ while maintaining time-invariant province specific mean-reversion relationships $g(\cdot)$ except for additive time-varying and province-specific macro shocks $\lambda_{p,t}$.

To understand the variation used to identify elasticities in this context, consider a setting with two provinces, where the one treated is identified by $R_i = 1$ (zero for the untreated province). We maintain the presence of two (four-year) differences ($t = 1, 2$). The reduced-form DDD regression takes the form

$$\begin{aligned}\Delta \log y_{it} = \lambda_P POST_t + \lambda_R R_i + \lambda_{PR}(POST_t \times R_i) + \mathbf{D}'_{it-4}\beta_D + (R_i \times \mathbf{D}_{it-4})'\beta_R \\ + (POST_t \times \mathbf{D}_{it-4})'\beta_P + (R_i \times POST_t \times \mathbf{D}_{it-4})'\theta + v_{it}.\end{aligned} \tag{9}$$

One can test whether $f(\cdot)$ is time-invariant by testing $\beta_P = 0$ or similarly whether $g(\cdot)$ is the same across provinces by testing $\beta_R = 0$. The coefficients θ retain their interpretation as behavioral responses in Jakobsen and Søgaard (2022) if they coincide with corresponding coefficients obtained from estimating the same DDD specification for $\log(1-\tau_{it})$ (or Δa_{it}). Plotting the coefficients θ gives a visual check on whether the maintained assumptions hold in the untreated region relative to the treated region of the (lagged) income distribution. The extension to the staggered design with multiple years is straightforward, replacing $POST$ with time fixed effects and multiple treated (and untreated) jurisdictions (provinces).

The estimating equation to obtain elasticities in (6) can be extended to the setting with multiple provinces. We have,

$$\Delta \log y_{it} = \alpha + \epsilon^c \Delta \log(1-\tau_{it}) + \gamma \Delta a_{it} + \mathbf{D}'_{it-4}(\beta_p + \beta_t) + \lambda_{t,p} + \nu_{it} \tag{10}$$

where the bin dummy controls and now allowed to have province specific coefficients and time-specific coefficients (additive interactions). The instruments, adapted to a multiple jurisdiction setting are $\Delta \log(1-\tilde{\tau}_{it})) = \log(1-\tau_{p,t}(y_{it-4})) - \log(1-\tau_{p,t-4}(y_{it-4}))$ and $\Delta \tilde{a}_{it} = a_{p,t}(y_{it-4})) - a_{p,t-4}(y_{it-4})$.

The variation that identifies ϵ^c and γ now comes from changes over time in the tax function which are province specific and vary across the (lagged) income distribution. We have already shown that provinces experience different changes in top rates at different moments in time. But exploiting the provincial variation adds another type of identifying variation under the maintained control function assumption. Changes in top rates occur in provinces with different thresholds for top rates. Hence, this generates for two taxpayers with the same income, independent variation in averate tax rates for a given change in the top rate. Therefore, we can allow for heterogeneity in ϵ^c and γ, allowing them to vary across different higher income groups.

4.4 Implementation details and diagnostics

Choice of provinces and income bins. Our empirical strategy requires a control sample of taxpayers at each lagged income level who were never subject to a tax reform. Our control observations are taxpayers in the provinces of Manitoba and Saskatchewan, two provinces in western Canada that did not introduce new tax brackets or reform statutory marginal tax rates in the relevant years. We exclude altogether from the analysis taxpayers in the four Atlantic provinces, several of which undertook a sequence of tax-increasing and tax-decreasing reforms, making them inappropriate as either treatment or control provinces.[10] These four provinces together comprise less than four percent of taxpayers in the top tax brackets affected by reforms.

Bins that define the vector $\mathbf{D}_{it-4}$ are equally distanced bin on the range of $\log y_{it-4}$ (between 75 thousands and 1 million). We consider jumps of 0.1 on the log-scale for the size of bins. In total, we have 27 bins. In all specifications, we control for age using a cubic age polynomial.

[10]The provinces of Newfoundland and Labrador and New Brunswick have both raised and lowered top marginal tax rates in recent years. In 2010, Nova Scotia implemented a reform that increased marginal tax rates for some high earners, while reducing them for others. For simplicity, we exclude Prince Edward Island, an Atlantic province with population less than 150,000 that experienced no reforms in the period of analysis.

Tax calculations. Our estimates require a measure of the marginal and average tax rates on reported income facing each taxpayer in each province and year. We calculate tax rates using tax schedule combining provincial and federal tax rates and brackets. Let the tax function $T_{p,t}(y)$ be parameterized as $T(y,\tau_{p,t},\kappa_{p,t})$ where $\tau_{p,t}$ is a vector of marginal tax rates in each bracket and $\kappa_{p,t} = (K_{p,t,1},\ldots,K_{p,t,N})$ the vector of thresholds (minimum incomes) for each bracket. Each tax function is piece-wise linear. Given this, the marginal tax rate is

$$\tau_{p,t}(y) = T'_{p,t}(y) = \tau_{p,t,k}, \quad \text{if} \quad y \in [K_{p,t,k}, K_{p,t,k+1}).$$

Likewise, the average tax rate is

$$a_{p,t}(y) = \frac{T(y,\tau_{p,t},\kappa_{p,t})}{y}.$$

For the instruments, we use the same function applied to taxable income in year $t-4$. While variation in $\tau_{p,t}$ was documented above, variation in thresholds over time exists because thresholds are indexed over time. Our calculations account for this variation. We show below, when discussing identification, how this variation correlates with changes in take-home rates.

Mean-reversion and difference lengths. We use four-year differences to measure our changes $\Delta \log y_{it}$, $\Delta \log(1-\tau_{it})$, and so on. For the post-treatment periods in our analysis, we pool observed changes in the year of the reform and the following year. Pre-reform control periods are the corresponding four-year changes measured two and three years prior to the reform year. As suggested by Weber (2014), the use of four-year differences reduces the extent of mean reversion at high levels of lagged income and so the potential bias in IV estimates using instruments constructed from lagged incomes. To illustrate this point, Figure 2 shows the mean reversion pattern estimated for various lag length. We pool data of the six treatment and control provinces, utilizing data for the years 2002 to 2010 where no tax change occurs. We use the implied annualized rate of change of income to compare on the same chart the average annual percent change for different lag length. The Figure shows that income changes become less negatively slopped as the lag length is increased. This is what one would expect if mean-reversion explains the negative correlation between lagged levels and changes.

Macro shocks and mean-reversion. The identifying assumption is that absent changes in tax rates, there are no province specific changes in the mean-reversion pattern over time. To check this, we consider a placebo period where there is actually a recession at the national level. We consider the estimation of our stacked reduced form DDD specification when using the 4 year changes in income in 2006 to 2010. This setup spans the Great Recession and therefore provides a good test for our identifying assumption. We would expect the DDD interactions (θ coefficients) to be jointly zero. Figure 3 shows these coefficient estimates along with confidence intervals. These DDD effects are very close to zero

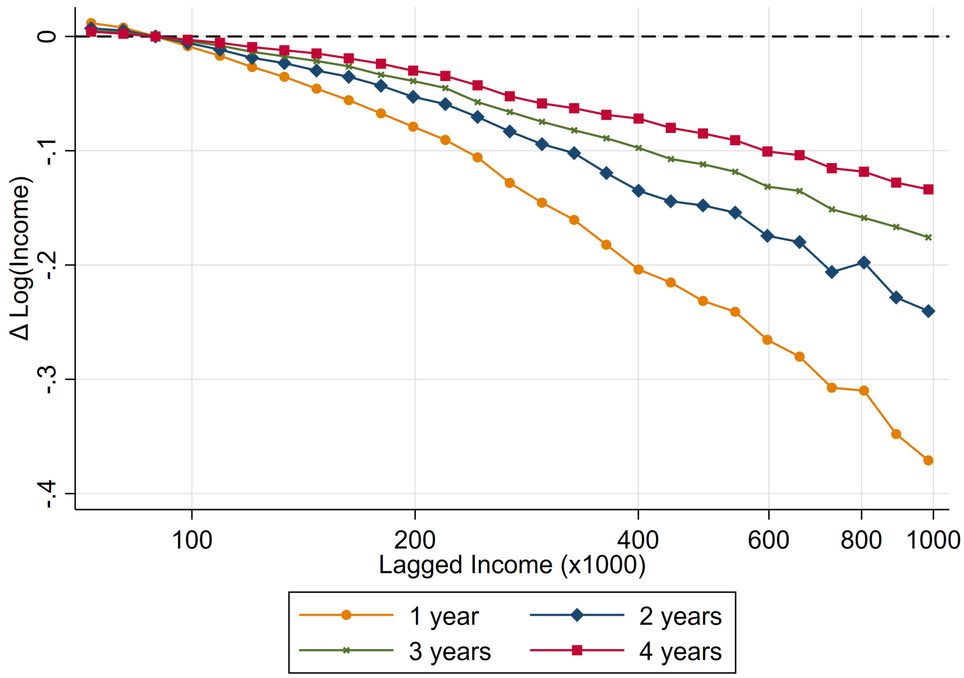

Figure 2: Patterns of Mean Reversion over Time. Shows estimates of a regression of log annualized changes in total income on bins of lagged (log) taxable income. Each line represents coefficients from a regression with a different lag length. Income changes of different lag lengths are annualized to make them comparable. Data are from the treatment and control provinces for the years 2002-10 where no tax change occurs.

and not statistically different from zero. The joint test statistic under the null that they are jointly zero has a relatively large p-value (p-value = 0.812). We also plot in the same figure the θ coeficients that would be obtained if we only considered the double-difference (differencing time and income level), therefore not exploiting the control provinces. In this case, we detect a (spurious) negative effect on reported lagged incomes at higher income levels (p-value < 0.001). This likely reflects the impact effect of the Great Recession on the distribution of top incomes. Hence, the difference-in-difference strategy is not robust to macro shocks in our setting.

Identification of income and substitution effects. Consider the variation in the instruments and how they identify income and substitution effects. First, we consider thresholds as fixed over time, $\kappa_{p,t} = \kappa_p$, and focus in particular on a change in the marginal tax rate above the top threshold for each province, which is denoted $K_{p,N}$. Then, the change in the simulated average tax rateis

$$\Delta a_{p,t}(y_{it-4}) = \Delta \tau_{p,t,N} \max\left(\frac{y_{it-4} - K_{p,N}}{y_{it-4}}, 0\right)$$

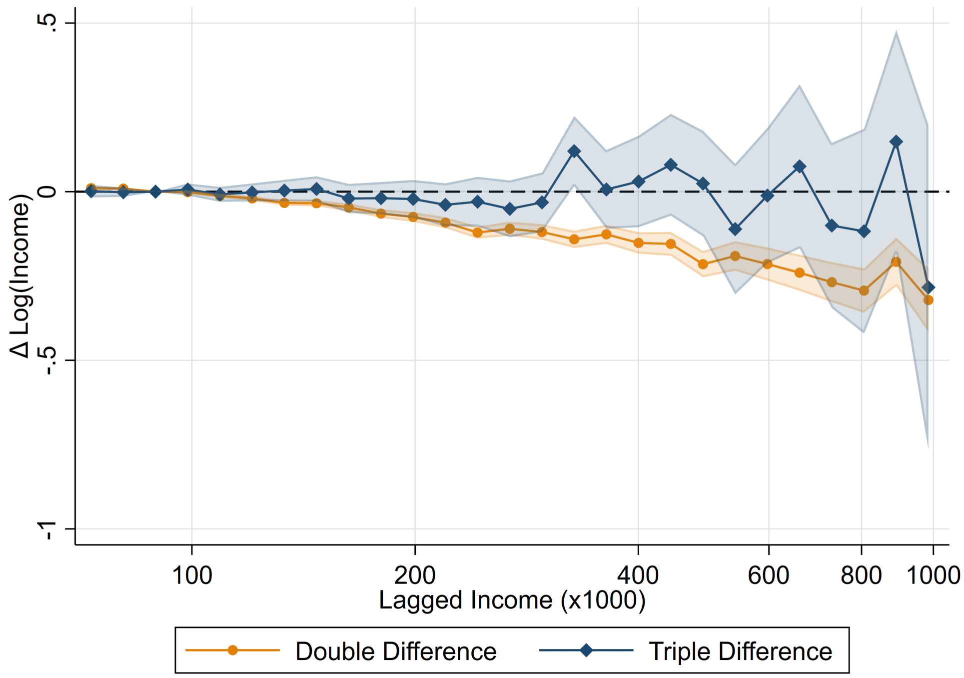

Figure 3: Placebo Test of DD and DDD Estimators. Reduced form DDD (triple difference) and DD (double difference) estimates and 95% confidence intervals. We define the pre period as 2006-07 and the post period to be 2009-2010. We compute 4-year changes for each period (i.e. 2006-2002 and 2007-2003, 2009-2005 and 2010-2006). All treated provinces are pooled together in these regressions. The DDD specification includes control provinces while the DD specification does not. The DD estimates likely reflect the impact effect of the Great Recession on the distribution of top incomes.

Given two reforming provinces, p and q where $K_{p,N} \leq K_{q,N}$ and a given income $y_{it-4} > K_{q,N}$, we can decompose the sources of variation as:

$$\Delta a_{p,t}(y_{it-4}) - \Delta a_{q,t}(y_{it-4}) = (\Delta\tau_{p,t,N} - \Delta\tau_{q,t,N})\left(\frac{y_{it-4} - K_{p,N}}{y_{it-4}}\right) + \Delta\tau_{q,t,N}\left(\frac{K_{q,N} - K_{p,N}}{y_{it-4}}\right)$$

The change in average tax rate comes differential variation in marginal tax rates but also for a given variation in marginal tax rates, from variation in thresholds. The larger the threshold difference $K_{q,N} - K_{p,N}$, the larger the difference in average tax rates for a given differential variation in marginal tax rates. Additional variation is provided across y_{it-4}, as is made clear from this decomposition. There is also variation coming from the fact that those with $K_{p,N} < y_{it-4} < K_{q,N}$ are treated in province p but not in province q but are potentially treated in province p. As we have shown, we have substantial differences in the thresholds at which tax rates change, which range from a low of \$100,000 in Quebec to \$300,000 for the largest tax increase in Alberta.

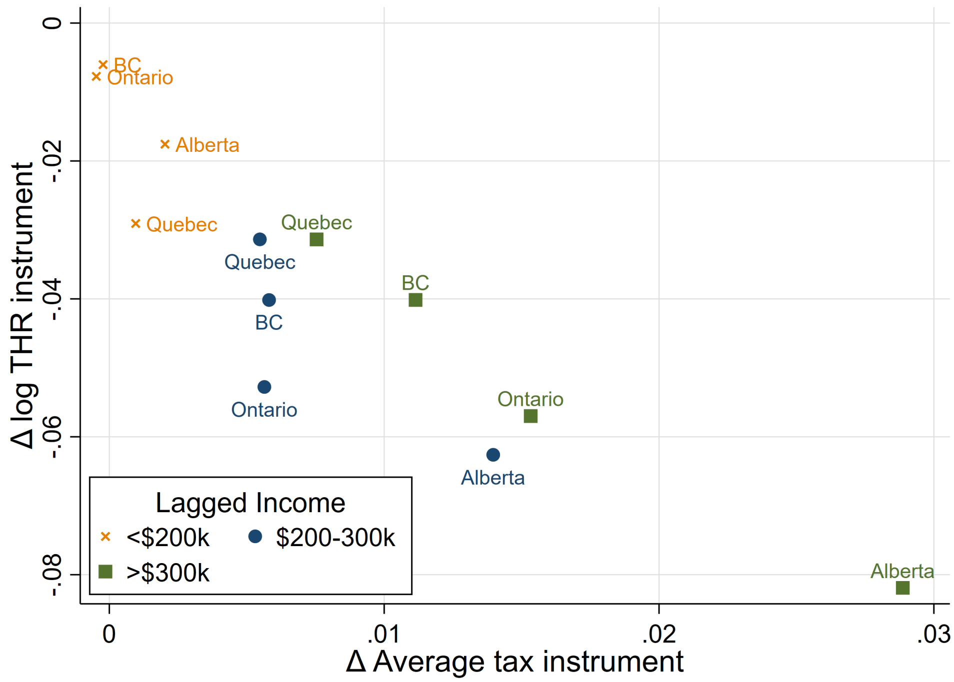

Figure 4: Variation in the Instruments. Each point in the figure plots the average simulated log take-home rate change $\Delta \log(1-\tau_{pt})$ against the simulated change in average tax rate change $(\Delta a_{p,t})$, for taxpayers in a particular province and a particular income range. Taxpayers are binned by four-year lagged income of $100,000-199,999, $200,000-299,999 and $300,000 and above. THR refers to take-home rate, BC to British Columbia.

Figure 4 plots for three lagged income groups the change in the log take-home rate instrument against the change in the average tax rate instrument. We do this for each of the treated provinces. There is generally a negative correlation between the instruments but substantial independent variation. Marginal tax rate changes rise with income in provinces that introduced multiple new top brackets. Average tax rate changes rise more sharply with income than do marginal tax rate changes, because of the mechanical effect of the reform on average tax rates.

5 Results

5.1 Reduced-form Evidence

We begin by presenting visual reduced-form evidence of the impact of tax changes in Figure 5. This figure displays estimates from Equation (10) for the log changes in income, as well as changes in the log take-home rate and the average tax rate. These regressions are conducted by combining all treated and control provinces.

(a) Change in Tax Measures

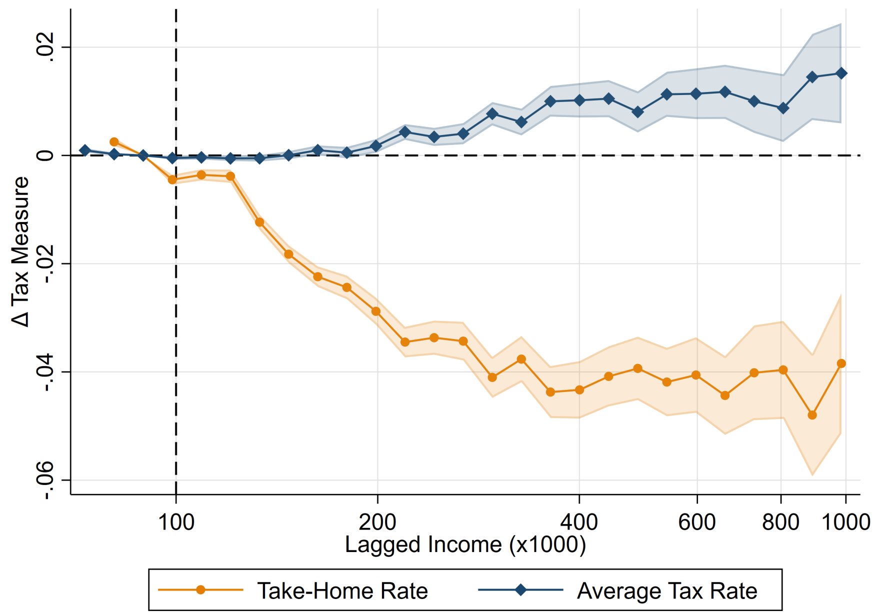

(b) Change in Log Total Income

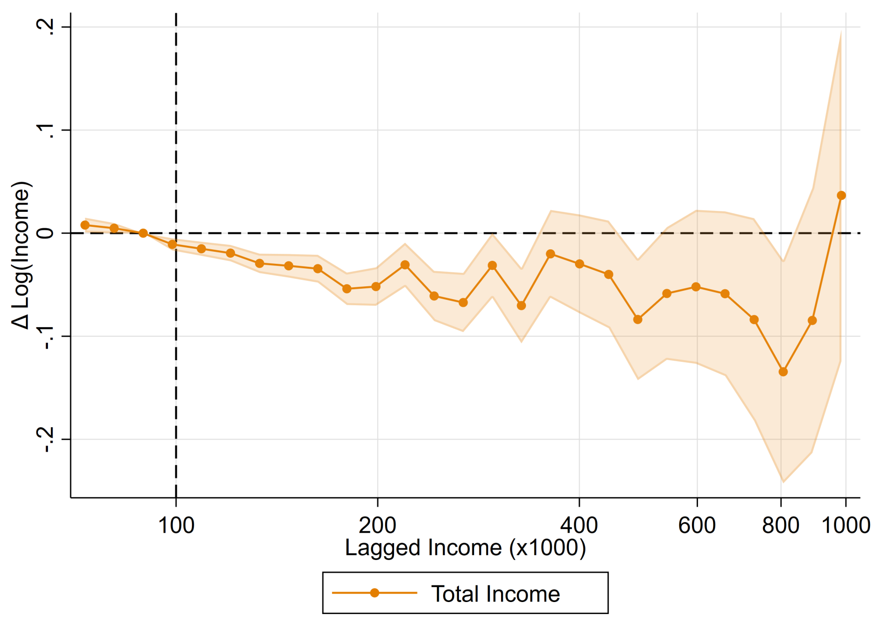

Figure 5: Triple Difference Estimates. The figures show estimates from the stacked triple difference specification, Equation (10). Panel (a) shows estimated triple-difference changes in the log of the take-home rate and in the average takes are due to tax reforms, and the vertical line represents the threshold at which the first reform applies (in Quebec at 100k). Panel (b) shows estimated triple-difference changes in total income.

First, panel (a) indicates that there is no change in the take-home rate or the average tax rate for incomes below \$100,000, which represents our untreated group. However, there is a statistically significant decline in the log take-home pay for incomes above this threshold in the treated provinces compared to the control provinces. This decline is accompanied by a corresponding increase in the average tax rate, relative to the control provinces. In panel (b), we observe a significant decrease in reported total income in the treated provinces when compared to the control provinces, with statistical significance up to approximately \$375,000. Although the decrease persists beyond this level, it is less precisely estimated due to sparse data at higher income levels. This suggests a positive elasticity of taxable income. The ratio of the coefficients from panel (b) to those from panel (a) provides an estimate of the elasticity of taxable income, which converges to 1.39 for incomes above \$200,000.[11]

Figure A.1 examines the robustness of the reduced-form results concerning our definition of the income base. We compare the response of total income (our preferred measure) with that of taxable income and T4 income. The results indicate that the responses across these tax measures are very similar. Additionally, we evaluate the robustness of these findings by including industry fixed effects to account for potential differences in industrial composition across provinces. Figure A.2 in the appendix presents estimates that closely resemble those reported in our baseline specification in Figure 5.

To explore how we might identify income and substitution effects, we next present the difference-in-difference-in-differences (DDD) specifications separately for each province. Due to smaller sample sizes at higher income levels, we focus on estimates for total incomes below \$300,000. Figure 6 displays the estimates, with panel (a) showing the change in the log of the take-home rate, while panel (b) illustrates the change in the log of total income. Notably, we observe that changes in the log take-home rate align well with the identified thresholds.

The reduced-form evidence shows a clear behavioral response to tax reforms in treated provinces. In panel (a), the largest changes in take-home rates are observed in Alberta, followed by Ontario, BC, and Quebec. If we focus solely on substitution effects, one would expect panel (b) to show the largest response in Alberta, followed by Ontario, BC, and finally Quebec. However, the largest response is actually observed in Quebec, while British Columbia shows the smallest response. This discrepancy may be attributed to the potential significance of income effects. At a given income level and for a specified change in the take-home rate, taxpayers in Quebec experience a more substantial drop in income than those in other provinces because they are the furthest from the threshold. Alberta exhibits a significant response due to a large change in the take-home rate. In contrast, British Columbia and Ontario have smaller responses because their thresholds are much higher. This distinction is crucial for differentiating between income effects and substitution effects.

[11] Figure 6 illustrates the estimated effects on income changes in each of the reforming provinces relative to the control group. We verify that the validation region holds for each province, accounting for different thresholds for the top bracket. The difference-in-differences pattern varies across provinces.

(a) Change in the log Take-Home Rate

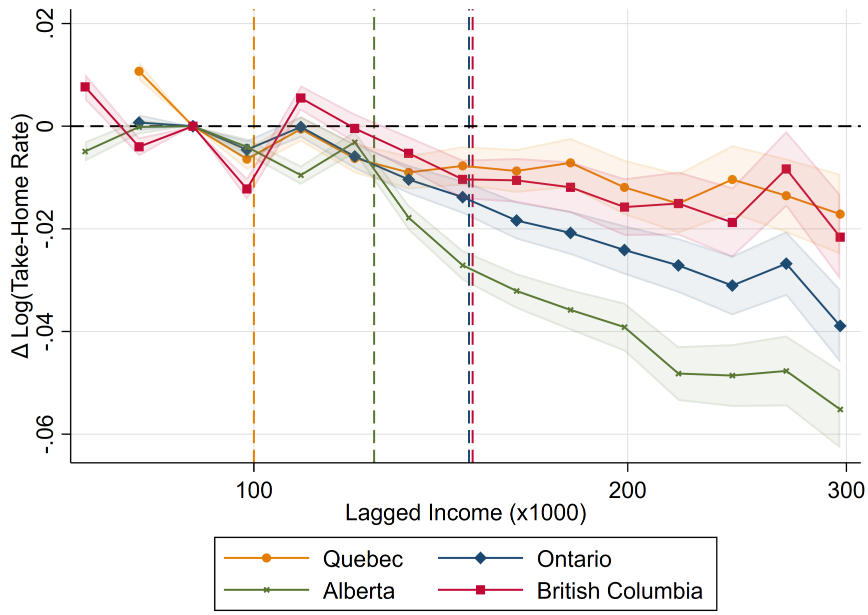

(b) Change in log Total Income

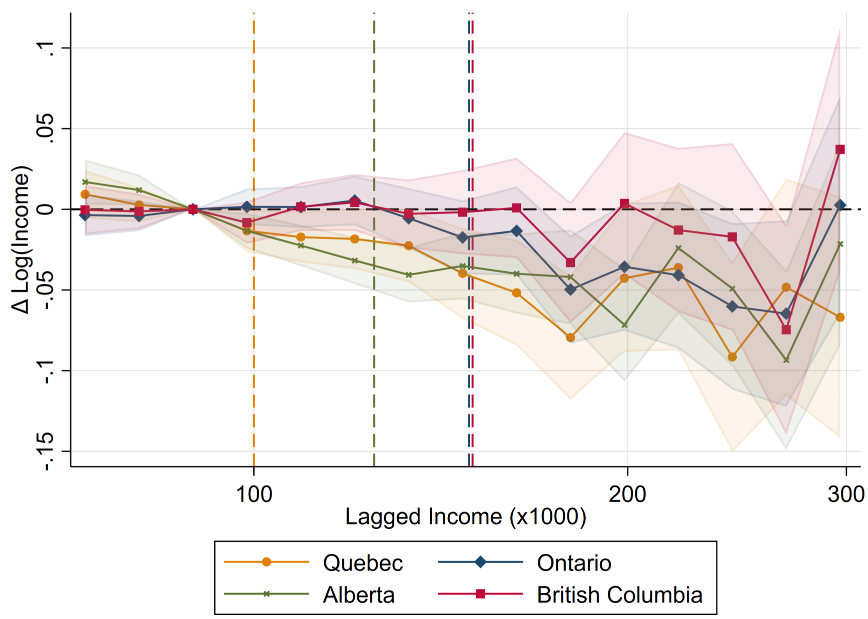

Figure 6: Triple Difference Estimates by Province. The figures show estimates from the triple difference specification by province. Panel A shows estimated triple-difference changes in the log of the take-home rate and vertical lines represent the thresholds at which reforms applied in each of the reforming provinces. Panel B shows estimated triple-difference changes in total income.

5.2 Elasticity estimates

We quantify substitution and income effects by estimating equation (10) using simulated take-home and average tax rates as instruments. Table 3 reports estimates.

When estimating a single elasticity of taxable income (ETI), we find an elasticity of 0.151 (se = 0.021), which falls within the lower range noted in the literature (column 1). Saez et al. (2012) suggests that a reasonable range for the ETI is between 0.2 and 0.4. Evidence from the reduced-form estimates indicates heterogeneity in responses. To investigate this heterogeneity, we interact the log of the take-home rate with indicators for whether lagged income is above \$200,000 and above \$300,000. Consequently, the elasticity estimate at higher income levels is the sum of the base elasticity (not interacted) plus the coefficients on the interactions. The estimates in column 2 reveal substantial heterogeneity. While the base elasticity is low for income just above the threshold (0.1), it rises significantly to 0.936 for incomes above \$200,000. Notably, we do not observe additional heterogeneity for incomes exceeding \$300,000. This pattern of increasing ETIs is also reported by Jakobsen and Søgaard (2022).

An alternative interpretation of the increasing elasticity pattern is that there is an income effect associated with changes in the average tax rate, which becomes more pronounced at higher income levels. We have shown theoretically that such negative effect are possible either due to salience effects or because of tax avoidance efforts which become more beneficial as the income loss increases. Column 3 demonstrates that we can detect this income effect, with an estimate of -3.658 indicating that a 1 percentage point change in the average tax rate results in a 3.658% decrease in reported income. In contrast, the average substitution effect is even larger, exhibiting an elasticity of 0.214. Due to variation in the thresholds across treated provinces, we can identify the income effect separately from heterogeneous substitution effects, as reported in column 4. Notably, the significant increase in elasticity beyond \$200,000 is partly attributable to the income effect, as the substitution elasticity decreases from 0.936 (column 2) to 0.743 (column 4). While the income effect remains important, its estimate is slightly lower, with reported income declining by 2.383% for every percentage point change in the average tax rate. When accounting for heterogeneity in both the substitution and income effects, the substitution elasticity beyond \$200,000 slightly decreases to 0.632. However, we find no evidence of heterogeneity in the income effect. Overall, we find that the increasing heterogeneity pattern in column 2 is a mix of heterogeneity in the substitution effect and an income effect.

We can assess sentivity of the elasticity estimates to various other modelling choices. Given the absence of heterogeneity in income effects, we focus on column 4 to report sensitivity. Results from various specifications are reported in Table 4.

We replicate in column 1 the estimates from our baseline specification (column 4) of Table 3. First, we consider using taxable income instead of total income (column 3). Consistent with what we observed with reduced-form evidence, estimates using taxable income are very similar to those with total income. This suggest little response beyond the reporting of total income. we assess how elasticity estimates vary when we control

Total Income	(1)	(2)	(3)	(4)	(5)
$\log(1-\tau_{p,t})$	0.151***	0.103***	0.214***	0.164***	0.169***
	(0.021)	(0.020)	(0.027)	(0.027)	(0.023)
$\log(1-\tau_{p,t}) \times I(y_{i,t-4} > 200k)$		0.833***		0.579**	0.463**
		(0.167)		(0.217)	(0.164)
$\log(1-\tau_{p,t}) \times I(y_{i,t-4} > 300k)$		0.071		−0.152	−0.052
		(0.206)		(0.240)	(0.217)
$a_{i,t}$			−3.658***	−2.382**	−2.510***
			(0.591)	(0.781)	(0.628)
$a_{i,t} \times I(y_{i,t-4} > 200k)$					0.354
					(0.304)
$a_{i,t} \times I(y_{i,t-4} > 300k)$					−0.420
					(0.465)
Number of observations	7 612 155	7 612 155	7 612 155	7 612 155	7 612 155
Instruments F-test	64 781	964	458	151	124

Table 3: 2SLS Estimates of Behavioral Elasticities. Each column reports an estimate of (10) for an alternative heterogeneity specification. Heterogeneity groups are defined as intervals of lagged income above \$100,000, \$200,000 and \$300,000. Columns 1-2 restrict income effects to be zero, and column 1 additionally restricts the take-home rate coefficient to be homogeneous. Columns 3-5 add income effects and successively greater heterogeneity in coefficients. All regressions include a 2-degree polynomial of age and fixed effects for income bin×year, bis×province and province×year. The last row of the table reports the Kleibergen-Paap rk Wald F-statistic of the instruments in the first stage. The average tax rate of our sample is 30% for lagged taxable incomes below 200k, 35% in the range of 200-300k and 38.5% for above. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

	(1)	(2)	(3)	(4)
$\log(1-\tau_{p,t})$	0.164***	0.170***	0.141***	0.198***
	(0.027)	(0.027)	(0.023)	(0.024)
$\log(1-\tau_{p,t}) \times I(y_{i,t-k} > 200k)$	0.579**	0.680**	0.662***	0.647***
	(0.217)	(0.220)	(0.192)	(0.191)
$\log(1-\tau_{p,t}) \times I(y_{i,t-k} > 300k)$	−0.152	−0.093	0.074	0.031
	(0.240)	(0.245)	(0.217)	(0.212)
$a_{i,t}$	−2.382**	−2.140**	−1.608*	−4.116***
	(0.781)	(0.790)	(0.645)	(0.737)
Number of observations	7 612 155	7 612 155	7 612 155	8 527 595
Income Measure	Total Income	Taxable Income	Total Income	Total Income
Sector-bins FE	No	No	Yes	No
Lag	4	4	4	3
Instruments F-test	151	151	197	200

Table 4: 2SLS Estimation Results with Alternative Specifications. Reports robustness of results to different specifications. Column 1 repeats our baseline results from column 4 of Table 3, where we include heterogeneity in the elasticity and an income effect. Column 2 uses taxable income in place of broad income as the dependent variable. Column 3 reports the baseline specification with additional controls for the interaction of income bins and industry (1-digit NAICS). Column 4 reports results with variables constructed on a lag of 3 years rather than 4. The last row of the table reports the KP rk Wald F-statistic of the instruments in the first stage. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

for industry fixed effects (column 2). Including industry fixed effect does not impact the substitution effect much but leads to a much smaller income effect (-1.608 instead of -2.382). Finally, we check whether using a 3-year lag instead of 4 changes substantially the estimates. The income effect increases while the substitution effect pattern remains similar. Overall, the pattern of substitution and income effects appears to robust to a number of modelling choices.

6 Policy implications

Our empirical results show that much of taxpayer responses to top bracket tax increases reflect the income effects of the reforms rather than substitution effects. This distinction has implications for the efficiency costs of income taxation, as well as the direction of welfare-improving reforms to the tax system.

Given the existing tax system $T(y)$, consider an increase in marginal tax rates in a small interval around an arbitrary income level y. This raises the average tax rate for all higher income levels, and it has behavioral effects to the resulting change in marginal and average

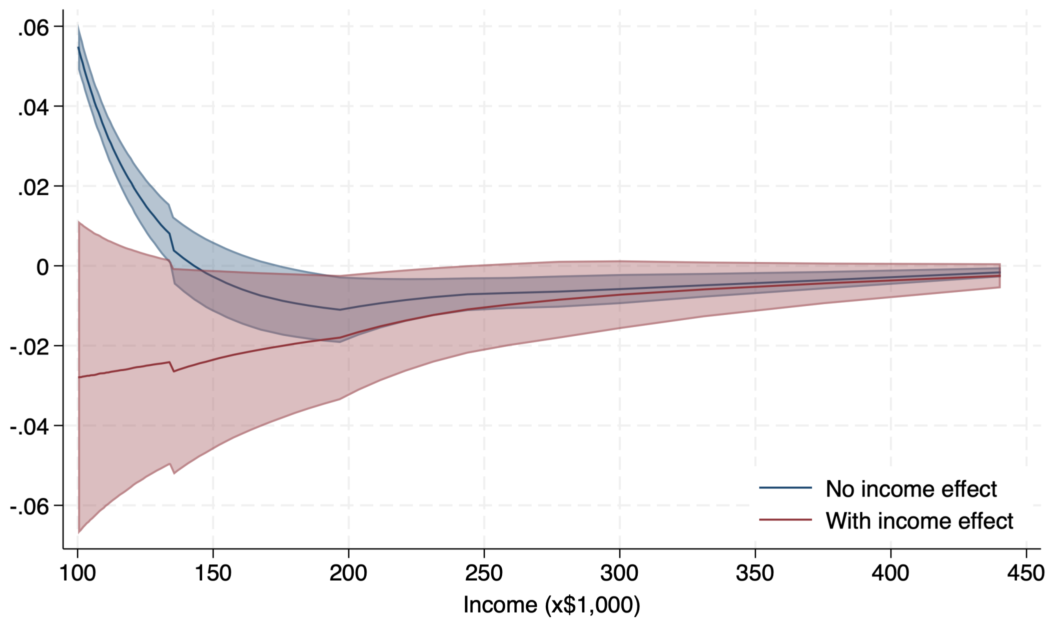

Figure 7: Gateaux derivatives of the revenue function. Depicted are the marginal revenue implications of a marginal increase in marginal tax rates in the neighborhood of each income level, computed from the formula in (11). Negative marginal revenue indicates that local tax reductions are Pareto-improving. When income effects are ignored, marginal revenue decreases sharply with income, which supports progressivity-decreasing reforms that decrease the top bracket rate while lowering the bracket threshold. When income effects are included, marginal revenue is relatively stable, which supports uniform rate reductions.

tax rates. The resulting change in revenues is proportional to the Gateaux derivative

$$dR(y) = (1 - F(y)) - \frac{T'(y)}{1 - T'(y)}\epsilon^c(y) y f(y) - \int_y \frac{T'(x)}{1 - T'(x)}\eta(x) dF(x) \tag{11}$$

The reform creates substitution effects on incomes locally and, through its effect on the average tax rate, it creates income effects at all higher levels of income. As emphasized by Bierbrauer et al. (2023), we can use the Gateaux derivatives to analyze directions of desirable reforms in marginal tax rates.

We simulate (11) using data on the milliciles of the Canadian income distribution in the year 2012 for incomes between \$100,000 and \$450,000, the range for which are elasticity estimates are valid. Using the quantile data, we estimate a smooth density of incomes $f(y)$ using the lowess filter. We simulate marginal revenues using the weighted average of federal and provincial marginal tax rates in that year, which were 43.4 percent for incomes up to \$135,000, and 46.4 percent at higher incomes. We calculate the estimated revenue effects using both the elasticities reported in column 2 of Table 3, where income effects are constrained to zero, and also for the general model with income effects reported in column 5 of Table 3. These heterogeneous elasticities were estimated for \$100,000 intervals of incomes; we linearly interpolate elasticities at each quantile of the data at the lower bounds of each heterogeneity group.

Figure 7 shows the estimated Gateaux derivatives, representing the revenue impacts of reforms that increase marginal tax rates in the neighborhood of each income quantile in our data. We estimate 95 percent confidence intervals of these estimates by simulating 10,000 draws from a joint normal distribution of the estimated elasticities, given the estimated covariance matrix of the estimates.

We find that marginal revenue from these reforms is negative in most cases. That is, the current system is judged to be on the "wrong side of the Laffer curve," so that local reductions in marginal tax rates would increase revenue and be Pareto-improving. This likely reflects the fact that our elasticity estimates are for unilateral changes in provincial tax rates in an environment where cross-provincial mobility of taxpayers and taxable income is high. A coordinated national tax increase might have a positive impact on revenue, but this effect cannot be estimated with our research design. The income tax base in Canada is co-occupied by provincial and the federal governments, leading to a vertical fiscal externality, in which the behavioral cost of tax increases by one level of government are partly borne by the revenues of the other level of government. As discussed in Milligan and Smart (2019), shared occupancy supports higher rates of overall redistributive taxation, in spite of the horizontal mobility of tax bases across provinces. Here, we abstract from the implications of estimates for fiscal federalism and focus on the implications of our estimates for the progressivity of the combined federal-provincial tax system.

That said, there are substantial differences in the results depending on whether income effects are included or not. Using the constrained estimates of column 2 without income effects, marginal revenue is positive for incomes below $145,000 (about the 97th percentile of the income distribution). Thus a reform that lowered the top-bracket threshold (which was then about $135,000) while lowering the top-bracket marginal tax rate would increase revenue. In this sense, the constrained estimates provide evidence to support reforms that *decrease* marginal rate progressivity at the very top of the income distribution. This reflects the much larger estimated substitution elasticities at higher income levels.

In contrast, using the unconstrained estimates of column 5, estimated marginal revenue negative and relatively stable for all incomes above $100,000. This reflects the relatively constant substitution elasticities over income levels in column 5, and that rate increases at lower income levels would increase average tax rates above, resulting in income effects on revenue that would offset much of the revenue gains. Thus the estimates of column 5 support a more progressive rate structure for top incomes than do the estimates of column 2, which ignore income effects of tax reforms.

7 Concluding remarks

We present new estimates of the compensated ETI based on administrative data around top-bracket tax reforms in Canada and a novel triple-difference estimation strategy. Our method arguably provides better identification than traditional approaches in the presence of temporal shocks to the income distribution, and it permits identification of both substitution and income effects of reforms when different jurisdictions change tax rates beginning at different income thresholds.

With few exceptions, the previous literature has estimated the ETI at different levels of lagged income while constraining the income effects of reforms to be zero. For top-bracket tax reforms, the magnitude of the income effect is an increasing function of a taxpayer's distance from the threshold at which the reform applies. Thus the income-related heterogeneity in the estimated ETI is confounded with the omitted income effect in much of the previous literature.

Our results indicate a strong negative effect of average tax rate changes on reported incomes, and so a positive income effect. We showed that such a response may be rationalized as an income effect on tax avoidance effort, or as a consequence of greater taxpayer attention to more salient tax reforms. Accounting for income effects reduces the income-related heterogeneity in the compensated ETI substantially but does not eliminate it entirely.

References

Bierbrauer, Felix J, Pierre C Boyer, and Emanuel Hansen, "Pareto-Improving Tax Reforms and the Earned Income Tax Credit," *Econometrica*, 2023, *91* (3), 1077–1103.

Burns, Sarah K and James P Ziliak, "Identifying the elasticity of taxable income," *Economic Journal*, 2016, *127* (600), 297–329.

Chetty, Raj, "Is the taxable income elasticity sufficient to calculate deadweight loss? The implications of evasion and avoidance," *American Economic Journal: Economic Policy*, 2009, *1* (2), 31–52.

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

Craig, Ashley C and Joel Slemrod, "Tax knowledge and tax manipulation: A unifying model," *Journal of Political Economy: Microeconomics*, 2024, *2* (2), 298–334.

Doerrenberg, Philipp, Andreas Peichl, and Sebastian Siegloch, "The elasticity of taxable income in the presence of deduction possibilities," *Journal of Public Economics*, 2017, *151*, 41–55.

Feldman, Naomi E, Peter Katuscak, and Laura Kawano, "Taxpayer confusion: Evidence from the child tax credit," *American Economic Review*, 2016, *106* (3), 807–835.

Feldstein, Martin, "Tax avoidance and the deadweight loss of the income tax," *Review of Economics and Statistics*, 1999, *81* (4), 674–680.

Gruber, Jon and Emmanuel Saez, "The elasticity of taxable income: evidence and implications," *Journal of Public Economics*, 2002, *84* (1), 1–32.

Jakobsen, Katrine Marie and Jakob Egholt Søgaard, "Identifying behavioral responses to tax reforms: New insights and a new approach," *Journal of Public Economics*, 2022, *212*, 104691.

Kleven, Henrik Jacobsen and Esben Anton Schultz, "Estimating taxable income responses using Danish tax reforms," *American Economic Journal: Economic Policy*, 2014, *6* (4), 271–301.

Kumar, Anil and Che-Yuan Liang, "Estimating taxable income responses with elasticity heterogeneity," *Journal of Public Economics*, 2020, *188*, 104209.

Milligan, Kevin and Michael Smart, "Taxation and top incomes in Canada," *Canadian Journal of Economics/Revue canadienne d'économique*, 2015, *48* (2), 655–681.

_ **and** _ , "An estimable model of income redistribution in a federation: Musgrave meets Oates," *American Economic Journal: Economic Policy*, 2019, *11* (1), 406–34.

Rees-Jones, Alex and Dmitry Taubinsky, "Measuring 'schmeduling'," *Review of Economic Studies*, 2020, *87* (5), 2399–2438.

Saez, Emmanuel, "Using elasticities to derive optimal income tax rates," *Review of Economic Studies*, 2001, *68* (1), 205–229.

_ **, Joel Slemrod, and Seth H Giertz**, "The elasticity of taxable income with respect to marginal tax rates: A critical review," *Journal of Economic Literature*, 2012, *50* (1), 3–50.

Slemrod, Joel and Wojciech Kopczuk, "The optimal elasticity of taxable income," *Journal of Public Economics*, 2002, *84* (1), 91–112.

Weber, Caroline E, "Toward obtaining a consistent estimate of the elasticity of taxable income using difference-in-differences," *Journal of Public Economics*, 2014, *117*, 90–103.

A Additional Reduced-Form Evidence

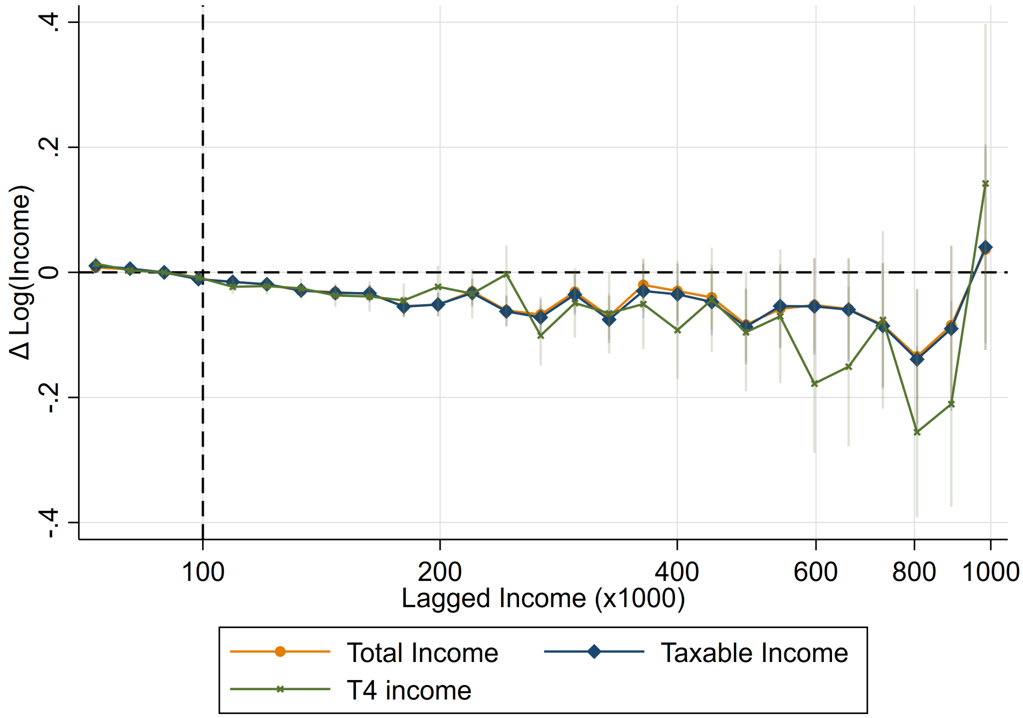

Figure A.1: Reduced-form estimates, by income measure. The figure shows estimates of the stacked triple-difference effects of reforms on successively narrower measures of the income base. Estimates correspond to those reported in Figure 5. Taxable income is Total income less certain deductions and exemptions. T4 income is employment income subject to third-party reporting to tax authorities.

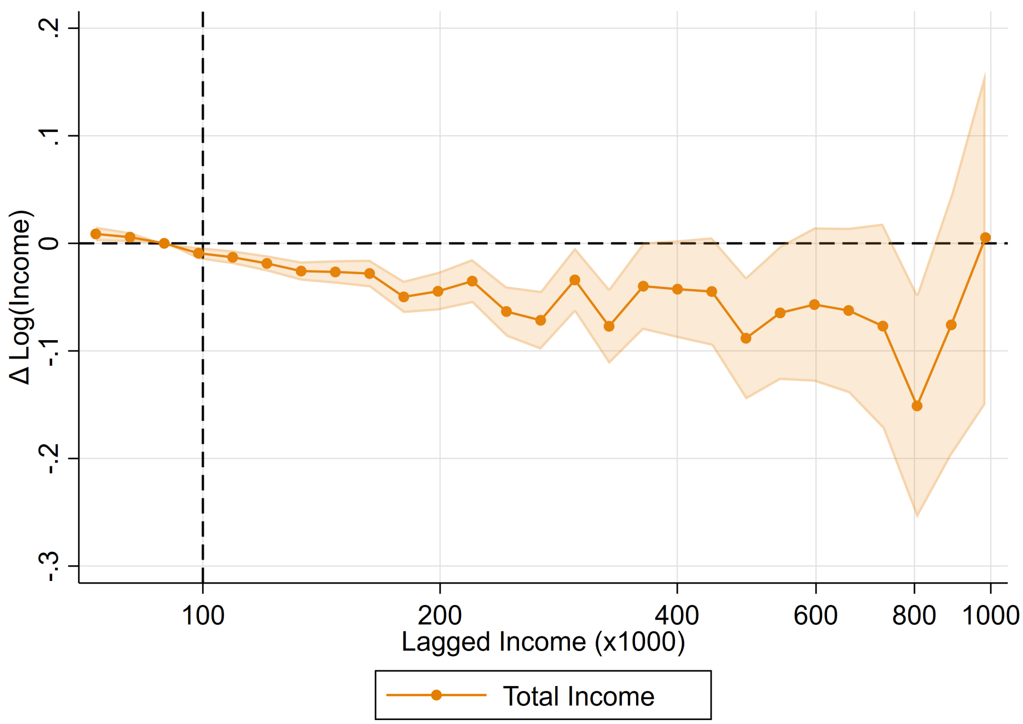

Figure A.2: Triple Difference Estimates Controlling for Industry Composition. The figure shows estimates of the stacked triple-difference effects of reforms accounting for industry * income bins fixed effects. Estimates correspond to those reported in Figure 5.

B Additional 2SLS Estimates

Table B.1: Estimation Results with Total Income, a Lag of 4 and industry FE

Total Income	(1)	(2)	(3)	(4)	(5)
$\log(1-\tau_{p,t}$	0.150***	0.098***	0.213***	0.141***	0.139***
	(0.020)	(0.019)	(0.025)	(0.023)	(0.020)
$\log(1-\tau_{p,t}) \times I(y_{i,t-4} > 200k$		0.837***		0.662***	0.604***
		(0.156)		(0.192)	(0.147)
$\log(1-\tau_{p,t}) \times I(y_{i,t-4} > 300k)$		0.211		0.074	0.134
		(0.196)		(0.217)	(0.198)
$a_{i,t}$			−3.471***	−1.608*	−1.517**
			(0.535)	(0.645)	(0.510)
$a_{i,t} \times I(y_{i,t-4} > 200k)$					0.107
					(0.278)
$a_{i,t} \times I(y_{i,t-4} > 300k)$					−0.434
					(0.433)
Number of observations	7 612 155	7 612 155	7 612 155	7 612 155	7 612 155
Sector-bins FE	Yes	Yes	Yes	Yes	Yes
Instruments F-test	70 637	1 072	543	197	162

Notes: The first column estimates baseline elasticity. The second column allows heterogeneity in elasticity along the income distribution. Coefficients at 200k and 300k report incremental effects. Column 3 reports the elasticity when introducing the income effect. Column 4 allows for heterogeneity in the elasticity, while column 5 also allows for heterogeneity in the income effect. Standard errors are clustered at the individual level and reported in parentheses. All regressions include a 2-degree polynomial of age and income bins-year, bins-province, province-year and industry-bins fixed effects. The last row of the table reports the KP rk Wald F-statistic of the instruments in the first stage.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

Table B.2: Estimation Results with Total Income and a Lag of 3

Total Income	(1)	(2)	(3)	(4)	(5)
$\log(1-\tau_{p,t})$	0.168***	0.103***	0.266***	0.198***	0.198***
	(0.018)	(0.017)	(0.026)	(0.024)	(0.021)
$\log(1-\tau_{p,t}) \times I(y_{i,t-3} > 200k)$		0.993***		0.647***	0.580***
		(0.142)		(0.191)	(0.151)
$\log(1-\tau_{p,t}) \times I(y_{i,t-3} > 300k$		0.328		0.031	0.087
		(0.172)		(0.212)	(0.200)
$a_{i,t}$			−5.983***	−4.116***	−4.094***
			(0.644)	(0.737)	(0.621)
$a_{i,t} \times I(y_{i,t-3} > 200k)$					0.162
					(0.326)
$a_{i,t} \times I(y_{i,t-3} > 300k)$					−0.430
					(0.634)
Number of observations	8 527 595	8 527 595	8 527 595	8 527 595	8 527 595
Instruments F-test	86 048	1 354	480	200	152

Notes: The first column estimates baseline elasticity. The second column allows heterogeneity in elasticity along the income distribution. Coefficients at 200k and 300k report incremental effects. Column 3 reports the elasticity when introducing the income effect. Column 4 allows for heterogeneity in the elasticity, while column 5 also allows for heterogeneity in the income effect. Standard errors are clustered at the individual level and reported in parentheses. All regressions include a 2-degree polynomial of age and income bins-year, bins-province and province-year fixed effects. The last row of the table reports the KP rk Wald F-statistic of the instruments in the first stage.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

Table B.3: Estimation Results with Taxable Income and a Lag of 4

Taxable Income	(1)	(2)	(3)	(4)	(5)
$\log(1-\tau_{p,t})$	0.168***	0.116***	0.234***	0.170***	0.176***
	(0.021)	(0.021)	(0.027)	(0.027)	(0.023)
$\log(1-\tau_{p,t}) \times I(y_{i,t-4} > 200k)$		0.909***		0.680**	0.564***
		(0.172)		(0.220)	(0.167)
$\log(1-\tau_{p,t}) \times I(y_{i,t-4} > 300k)$		0.107		−0.093	0.006
		(0.214)		(0.245)	(0.222)
$a_{i,t}$			−3.806***	−2.140**	−2.298***
			(0.609)	(0.790)	(0.635)
$a_{i,t} \times I(y_{i,t-4} > 200k)$					0.370
					(0.309)
$a_{i,t} \times I(y_{i,t-4} > 300k)$					−0.388
					(0.476)
Number of observations	7 612 155	7 612 155	7 612 155	7 612 155	7 612 155
Instruments F-test	64 781	964	458	151	125

Notes: The first column estimates baseline elasticity. The second column allows heterogeneity in elasticity along the income distribution. Coefficients at 200k and 300k report incremental effects. Column 3 reports the elasticity when introducing the income effect. Column 4 allows for heterogeneity in the elasticity, while column 5 also allows for heterogeneity in the income effect. Standard errors are clustered at the individual level and reported in parentheses. All regressions include a 2-degree polynomial of age and income bins-year, bins-province and province-year fixed effects. The last row of the table reports the KP rk Wald F-statistic of the instruments in the first stage.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

Table B.4: Estimation Results with T4 Income and a Lag of 4

T4 Income	(1)	(2)	(3)	(4)	(5)
$\log(1-\tau_{p,t})$	0.075*	0.022	0.130***	0.103**	0.085**
	(0.033)	(0.033)	(0.037)	(0.035)	(0.033)
$\log(1-\tau_{p,t}) \times I(y_{i,t-4} > 200k)$		0.729***		0.305	0.270
		(0.170)		(0.214)	(0.205)
$\log(1-\tau_{p,t}) \times I(y_{i,t-4} > 300k)$		0.376		−0.019	−0.171
		(0.212)		(0.243)	(0.268)
$a_{i,t}$			−4.945***	−4.175***	−3.367***
			(0.764)	(0.983)	(0.891)
$a_{i,t} \times I(y_{i,t-4} > 200k)$					−0.427
					(0.533)
$a_{p,t} \times I(y_{i,t-4} > 300k)$					−1.813
					(0.940)
Number of observations	6 529 070	6 529 070	6 529 070	6 529 070	6 529 070
Instruments F-test	72 856	22 046	612	205	153

Notes: The first column estimates baseline elasticity. The second column allows heterogeneity in elasticity along the income distribution. Coefficients at 200k and 300k report incremental effects. Column 3 reports the elasticity when introducing the income effect. Column 4 allows for heterogeneity in the elasticity, while column 5 also allows for heterogeneity in the income effect. Standard errors are clustered at the individual level and reported in parentheses. All regressions include a 2-degree polynomial of age and income bins-year, bins-province and province-year fixed effects. The last row of the table reports the KP rk Wald F-statistic of the instruments in the first stage.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.