MEMORANDUM

No 04/2022
October 2022

Behavioural responses to income taxation in Norway

**Michael Graber, Magne Mogstad,
Gaute Torsvik and Ola Vestad**

ISSN: 0809-8786

Department of Economics
University of Oslo

This series is published by the
University of Oslo
Department of Economics

P. O.Box 1095 Blindern
N-0317 OSLO Norway
Telephone: + 47 22855127
Fax: + 47 22855035
Internet: http://www.sv.uio.no/econ
e-mail: econdep@econ.uio.no

In co-operation with
The Frisch Centre for Economic Research

Gaustadalleén 21
N-0371 OSLO Norway
Telephone: +47 22 95 88 20
Fax: +47 22 95 88 25
Internet: http://www.frisch.uio.no
e-mail: frisch@frisch.uio.no

Last 10 Memoranda

No 03/2022	Eskil Heinesen, Christian Hvid, Lars Kirkebøen, Edwin Leuven and Magne Mogstad Instrumental variables with unordered treatments: Theory and evidence from returns to fields of study
No 02/2022	Mikkel Gandil and Edwin Leuven College admission as a screening and sorting device
No 01/2022	Askill H. Halse, Karen E. Hauge, Elisabeth T. Isaksen, Bjørn G. Johansen and Oddbjørn Raaum Local Incentives and Electric Vehicle Adoption
No 04/2021	Lars Kirkebøen School value-added and long-term student outcomes
No 03/2021	Mikkel Høst Gandil Substitution Effects in College Admissions
No 02/2021	Karine Nyborg A perfectly competitive economy is an economy without welfare relevant endogenous learning
No 01/2021	Elisabeth T. Isaksen & Bjørn G. Johansen Congestion pricing, air pollution, and individual-level behavioral responses
No 02/2020	Manudeep Bhuller, Tarjei Havnes, Jeremy McCauley & Magne Mogstad How the Internet Changed the Market for Print Media
No 01/2020	Anders Kjelsrud and Kristin Vikan Sjurgard Public work and private violence
No 10/2019	Manudeep Bhuller, Andreas R. Kostøl and Trond C. Vigtel How Broadband Internet Affects Labor Market Matching

Previous issues of the memo-series are available in a PDF® format at:http://www.sv.uio.no/econ/english/research/unpublished-works/working-papers/

Behavioral responses to income taxation in Norway

Michael Graber* Magne Mogstad† Gaute Torsvik‡ Ola L. Vestad§

Abstract

In this report, we combine theory and empirical estimates for how labor earnings respond to changes in tax rates and unearned income. We use lottery winnings to obtain variation in unearned income and tax reforms to obtain variation in the net of tax rate. Combining this information with measures of extensive margin responses and the progressivity of the Norwegian income tax schedule, we are able to point identify uncompensated and compensated behavioral responses to income taxes and therefore to calculate efficiency losses and optimal income tax rates (for given welfare weights).

JEL-codes: D15, J22, H21, H31, H53

Key words: income effect, labor supply elasticities, lottery winnings, efficiency loss, optimal income taxation.

*SSB

†University of Chicago

‡University of Oslo

§SSB

1 Introduction

The goal of this report is to estimate the behavioral responses to income taxation in Norway, and to explore the economic implications of these responses for tax policy. To achieve this goal, it is necessary to consider the three key economic criteria for assessing the impact of a tax or a tax reform: (i) how does it affect government revenue, (ii) how does it affect economic efficiency, and (iii) how does it affect the distribution of disposable income and welfare across households. What makes it difficult to answer these questions is that individuals will re-optimize and adjust their behavior when a tax is introduced or changed.

A change in the income tax may affect two important types of labor market decisions; whether or not to participate in the labor market and the hours of work (or effort) if participating. Participation will typically depend on the difference in the income disposable for consumption when working versus not working, as well as the pecuniary and non-pecuniary costs and benefits associated with working. Income taxation will play a role here because it affects the difference in disposable income if working versus not working. What matters then is the average taxes paid on earned income as well as benefits received if not working. By comparison, among those that participate in the labor market, it is the marginal income tax rate that matters for the choice of work hours (or effort). Economic theory tells us that the intensive margin response to a change in a price – such as a wage, the price on leisure – can be decomposed into a substitution and an income effect. The substitution effect captures the labor supply response to a change in the tax rate assuming taxpayers obtain a lump sum compensation to maintain the utility they had before the tax change. The income effect captures how labor market earnings responds to an increase in unearned income.

An evidence based assessment of the revenue, efficiency, and distributional effects of an income tax reform requires credible estimates of the relevant behavioral responses, both at the extensive and intensive margin. Economic theory provides limited guidance here, and it is challenging to use data to estimate the relevant behavioral responses. The revenue effect of a tax reform can be estimated without separating between extensive and intensive margin responses or between substitution and income effects at the intensive margin. All that is needed is to find exogenous variation in tax rates (or wages) that can be used to identify how changes in income taxation affect taxable income and government revenue. In contrast, calculating the efficiency loss of income taxation requires knowledge of the substitution effect of the tax change. The substitution effect determines the compensated earnings response to a tax change, which cannot be directly recovered from even a controlled trial that randomly varies wages or tax rates. Yet, credible estimates of the substitution effect is (or at least should be) essential for any conclusion drawn about the whether taxes should be increased or decreased.

In much of the empirical public finance literature, the difficulty of estimating the efficiency costs of income taxation is often circumvented simply by assuming away income effects. If income effects are assumed away, the compensated and uncompensated earnings responses are, by assumption, identical. Under this restriction, one may therefore use (exogenous) variation in tax rates alone to assess both revenue and efficiency effects of taxation. As discussed later, much of the existing empirical work reports fairly modest earnings responses to changes in income taxation, at least on the intensive margin. In the absence of income effects, these findings imply small substitution effects, and little efficiency loss from income taxation. However, in the presence of income effects, a small uncompensated response could be a result of the substitution and income effects both being large or both being small. For efficiency and welfare assessments, distinguishing between these two scenarios is essential.

Our report differs from much of the empirical public finance literature in one important way: we take the possibility of non-negligible income effects seriously. Instead of assuming no income effects, we will combine economic theory and data variation in both tax rates and unearned income to obtain estimates of uncompensated and compensated behavioral responses to income taxes. Armed with these estimates, we will analyze the deadweight loss of income taxation and optimal income taxation.

The report is organized as follows: In Section 2, we present a model of income taxation where the government uses a non-linear income tax to collect revenue from individuals with different labor market productivity and preferences who can make decisions at both the extensive and intensive margins. The model lays the groundwork for our empirical analysis, defining the target parameters. In Section 3, we use a series of Norwegian tax reforms to analyze how individuals' earnings respond to tax reforms. In Section 4, we study how individuals respond to idiosyncratic and exogenous changes in household wealth and unearned income. Our analyses combine administrative data on Norwegian lottery winners with an event-study design that exploits variation in the timing of lottery wins. In Section 5, we combine the model in Section 2 with the estimates from Sections 3 and 4 to recover income and substitution effects across the earnings distribution. We then analyze the deadweight loss of income taxation and optimal income taxation.

The main findings from our analyses may be summarized with five broad conclusions. First, we find significant but modest earnings responses to changes in income taxation. There is considerable heterogeneity across the income distribution, with the largest earnings responses in the lower part of the distribution. Second, the earnings responses to the windfall gains from lottery winnings suggest significant and sizable reductions in earnings along both the intensive and extensive margin. These effects are heterogeneous across the income distribution, with households in higher quartiles of the income distribution reducing their earnings by a larger amount. Third, income effects are considerable, especially for high income earners who reduce their earnings by 76 cents per dollar increase in unearned income. The compensated elasticities range from 0.28 for low-income earners to 0.15 for medium-income earners. The uncompensated elasticities are closer to zero, never exceeding 0.23 in absolute value. Fourth, we find that an increase in the top marginal tax rate would come at the cost of a considerable loss in efficiency. For an extra dollar of taxes raised, an increase in the top tax rate would impose an additional cost of around 59 cents on high earning individuals (on top of the dollar paid in taxes). Lastly, to justify the current level of the top income tax, the decision maker needs to be indifferent between one more dollar of public funds and around 1.16 dollars of additional disposable income to high income earners. If the marginal value of public funds is higher (lower), then the top income tax rate should be increased (decreased).

2 Theory of income taxation

In this section, we present a model of income taxation where the government uses a non-linear income tax to collect revenue from individuals with different labor market productivity and preferences and who make decisions at both the extensive and intensive margins. The model lays the groundwork for our empirical analysis, defining the target parameters for policy analysis.

2.1 Income taxation and taxable income

In the standard economic model of labor supply, individuals derive utility from consuming leisure and market goods and transform leisure to market goods by supplying work hours in the labor market. There is a large literature, both theoretical and empirical, that examines how income taxation affects hours of work.

Modern public finance has shifted focus away from how income taxation affects hours worked to its effect on taxable income. This is a more general approach since work hours is only one of many ways individuals may respond to changes in income taxation. They also choose how much effort to exert at work, how much to invest in training, which occupation to choose, and how much time, money and effort they should use to avoid income taxes. All of these choices may depend on income taxation.

Feldstein (1995) gave a formal and welfare based reason for when and why it is sufficient to measure the taxable income impact of all these different responses to income taxation. He showed that if workers optimize across these different margins, the only measures needed in order to assess both the efficiency, revenue, and welfare effects of a change in the income tax is the impact it has on compensated and uncompensated responses in taxable income. Later, Chetty (2009) showed that this result hinges on a few assumptions, for example that the costs one person pays in order to reduce taxable income does not become the income of another person (a lawyer).

Access to reliable data is another reason for estimating tax effects on taxable income rather than work hours. Taxable income is reported to the tax authorities and recorded in government registers. Data on taxable income therefore have larger sample sizes (often covering the entire working population) and fewer measurement errors than survey data on work hours. Furthermore, in survey data, wage rates are often imputed by dividing labor income on reported hours of work. Measurement errors in work hours will then lead to a division bias in the estimation of labor supply.

For these reasons, we will throughout this report focus on how taxable income is affected by income taxation. Since we focus on workers and labor income taxation, taxable income will be defined as the sum of wages and self-employment income. We will use the terms earnings and taxable income interchangeably.

Textbook model

Consider a simple static case with no heterogeneity and with a linear income tax. An individual chooses earnings y and consumption c to maximize

$$u(c, y) \tag{2.1}$$

subject to the intratemporal budget constraint

$$c = (1 - \tau)y + R. \tag{2.2}$$

The utility function u is a twice-differentiable and concave in $(c, -y)$. The first order condition to this problem, $u_c(1 - \tau) + u_y = 0$, implicitly defines optimal earnings as a function of the net of tax rate and unearned income: $y^* = y^*(1 - \tau, R)$.

We can use the Slutsky equation to decompose a change in the net of tax rate into a price (substitution) and an income effect. Let ζ^u denote the uncompensated or Marshallian net of tax elasticity $\zeta^u \equiv \frac{\partial y^*}{\partial(1-\tau)} \frac{(1-\tau)}{y^*}$ and $\eta \equiv \frac{\partial y^*}{\partial R}$ denote the income parameter (the effect of a small increase in unearned income on earned income). The substitution effect on earnings (the Hicksian earnings response) is given by the earnings response that minimizes the costs needed to reach a given utility level at different net of tax rates; $\zeta^c = \frac{\partial y^*}{\partial(1-\tau)} \frac{(1-\tau)}{y^*} |_{u=\overline{u}}$. The relationship between these parameters is given by the Slutsky equation

$$\zeta^u = \zeta^c + (1-\tau)\eta. \tag{2.3}$$

These parameters carry all the information needed to assess how changes in unearned income and changes in the tax rate affect government revenue and economic efficiency, if all earnings responses to taxation are at the intensive margin. This is, however, a restrictive setting. In order to make the model more applicable for practical policy analysis it should be extended in several directions. A policy relevant framework should allow for (i) individual heterogeneity, (ii) extensive margin decisions (participate in the labor market or not), (iii) more complex (non-linear) tax schemes, and (iv) dynamics (savings decisions).

2.2 A dynamic heterogeneous-agent economy with non-linear taxes and extensive margin decisions

The economy now consists of heterogeneous households. Each household i lives for $T \leq \infty$ periods and solves

$$\max_{\{c_{i,t}, y_{i,t}, n_{i,t}\}_t} \sum_{t=1}^{T} \beta^{t-1} \mathfrak{u}_i(c_{i,t}, y_{i,t}), \tag{2.4}$$

subject to a sequence of intra- and inter-temporal budget constraints,

$$c_{i,t} = y_{i,t} - \mathcal{T}(y_{i,t}) + n_{i,t},$$
$$n_{i,t} = (1+r)a_{i,t-1} - a_{i,t},$$

where $a_{i,0}$ are given initial assets of household i. $\mathcal{T}(\cdot)$ is a twice-differentiable tax function with $\mathcal{T}', \mathcal{T}'' \geq 0$, and r is the post-tax interest rate that satisfies $r = \beta^{-1} - 1$.

The utility function $\mathfrak{u}_i$ takes the form

$$\mathfrak{u}_i(c,y) = \begin{cases} u_i(c,y) & \text{if } y > 0 \\ u_i(c,y) + \xi_i & \text{if } y = 0 \end{cases}, \tag{2.5}$$

where u_i is a twice-differentiable function that is concave in $(c, -y)$, and $\xi_i \geq 0$. The parameter ξ_i captures utility from non-working and is a simple modeling device to capture extensive-margin responses (see, e.g., Kleven and Kreiner, 2006; Jacquet et al., 2013).

Individuals are heterogeneous in their initial wealth levels $a_{i,0}$, utility functions u_i, and benefits of non-working ξ_i. Hence, individuals have different optimal levels of earnings $\left\{y^*_{i,t}\right\}_{i,t}$. However, because the dynamic problem of the household is stationary, $y^*_{i,t}$ is independent of t. This feature of the model substantially simplifies our arguments and exposition, as we can drop subscripts t from the analysis. This model specification, however, is not restrictive for our analysis.[1]

Let $H(y)$ be the earnings distribution in any period. We assume that $H(y)$ has a well-defined density $h(y)$. For any variable $\{x_i\}_i$, we use $\mathbb{E}x_i$ to denote its population average in equilibrium given function $\mathcal{T}(\cdot)$. We use $\mathbb{E}_y$ and $\mathbb{E}_{\geq y}$ to denote the average over all households who have earnings $y_i = y$ and $y_i \geq y$, respectively. Aggregate earnings Y in this economy are

$$Y \equiv \mathbb{E}y^*_i = \int [\mathbb{E}_y y^*_i] h(y) dy.$$

We use a perturbational approach to study how individuals respond to changes in unearned income and tax rates. Since the model allows for extensive-margin responses, small perturbations of taxes and transfers may have a discontinuous effect on optimal choices of any household i. We assume that the underlying heterogeneity in preferences is such that the earnings density $h(y)$ changes smoothly to all perturbations we consider below.

Note that since the utility function is time-separable, household i's problem (2.4) can be split into two subproblems (see, e.g., MaCurdy, 1983; Blomquist, 1985; Blundell and Walker, 1986). An outer subproblem characterizes the optimal choice of unearned income $\left\{n^*_{i,t}\right\}_t$, where $n^*_{i,t}$ is independent of t given our assumption that $\beta^{-1} = 1 + r$. The optimal bundle (c^*_i, y^*_i) is determined by the trade-off between the intensive and extensive margins. The optimal intensive-margin response $(c^{\text{int}}_i, y^{\text{int}}_i)$ is the solution to

$$\max_{c_i, y_i} u_i(c_i, y_i) \tag{2.6}$$

subject to

$$c_i = y_i - \mathcal{T}(y_i) + n^*_i. \tag{2.7}$$

The choice between intensive and extensive-margin responses is then determined by

$$\max \left\{ u_i(c^{\text{int}}_i, y^{\text{int}}_i), u_i(n^*_i - \mathcal{T}(0), 0) + \xi_i \right\}.$$

Importantly, the same properties of u_i that are summarized by income and substitution effects in (2.3) also govern the optimal intensive margin responses to perturbations in problem (2.6). These parameters are therefore key to understand how income taxation influence economic efficiency, government revenue, and the distribution of welfare. However, since income taxation may have extensive margin labor market responses it is not necessarily sufficient to know the Slutsky parameters (defined in equation 2.3) in order to evaluate tax reforms. Furthermore, extensive margin responses make it more difficult to use reforms to identify the Slutsky parameters.

In the remainder of this section, we will show formally how changes in unearned income and marginal tax rates relate to the Slutsky parameters, and how these parameters are informative about tax policy. Later

[1] Golosov et al. (2021) solve the model with overlapping generations and show that the same results can be derived in such a framework.

we will combine these theoretical results with estimates of the earnings responses to tax reforms and lottery winnings to bound or point identify the Slutsky parameters and to draw inferences about tax policy.

Changes in unearned income

In a model with no extensive margin and with a linear income tax the earnings response to an exogenous change in unearned income can be used to identify the income effect η in the Slutsky equation. This is not the case if the government uses a non-linear income tax and individuals can choose to leave the labor market in response to receiving a windfall of unearned income.

Assume the wealth of household i increases unexpectedly by the amount L in period t that we set, without loss of generality, to 1. We are interested in the marginal impact of this windfall. For any equilibrium variable $x(L^{\text{ann}})$ that is differentiable in L^{ann}, let the marginal effect of this increase, $\partial_L x$, be defined as $\partial_L x \equiv \lim_{L^{\text{ann}}\to 0} \frac{x(L^{\text{ann}})-x(0)}{L^{\text{ann}}}$, where L^{ann} is the annuitized value of L. We choose this definition so that $\partial_L n_i^* = 1$ and all marginal responses are expressed per dollar increase in unearned income in a given period.

Since households may respond to a change in wealth along the extensive margin, individual responses, such as $\partial_L y_i^*$, may not be well defined. Instead, we focus on earnings responses averaged over groups of households. Let $\Lambda_L(y, L^{\text{ann}})$ be the fraction of households with earnings y who would stop working if their unearned income increased by L^{ann}. The parameter $\lambda_L(y) \equiv \partial_L \Lambda_L(y, L^{\text{ann}})$ is the marginal propensity to stop working for households with earnings y. The marginal propensity to earn out of unearned income (MPE) is defined as

$$\text{MPE}(y) \equiv \partial_L \mathbb{E}_y y_i^*.$$

The population averages of these objects are the averages of MPE(y) over the entire earnings distribution:

$$\text{MPE} \equiv \int \text{MPE}(y) dH(y).$$

The marginal propensity to stop working, $\lambda_L(y)$, is well defined because we assume that the earnings density responds smoothly to all perturbations. This, in turn, implies that the MPE(y) is well defined for all y and satisfies

$$\text{MPE}(y) = \text{MPE}^{\text{int}}(y) + \text{MPE}^{\text{ext}}(y). \tag{2.8}$$

Equation (2.8) shows that the expression for the MPE(y) consists of two terms, capturing intensive- and extensive-margin responses. The extensive-margin response is given by

$$\text{MPE}^{\text{ext}}(y) \equiv -y\lambda_L(y). \tag{2.9}$$

This response is proportional to $\lambda_L(y)$, which is simply the average marginal propensity to stop working of someone making y dollars. The intensive margin response is given by

$$\text{MPE}^{\text{int}}(y) \equiv \mathbb{E}_y \frac{\eta_i}{1+\theta_i \zeta_i^{\text{c}}}, \text{with } \theta_i \equiv \frac{\mathcal{T}''(y_i^*)y_i^*}{1-\mathcal{T}'(y_i^*)}. \tag{2.10}$$

The intensive-margin responses are proportional to $\mathbb{E}_y \frac{\eta_i}{1+\theta_i \zeta_i^c}$. To understand why they take this form, observe that an increase in unearned income has two effects. The direct effect is the standard income response that calls for a reduction in earnings by η_i. The second effect is indirect, and is driven by the non-linearity in the tax function $\mathcal{T}(\cdot)$. When a household reduces its earnings, the marginal tax rate that it faces may fall due to tax progressivity. The substitution effect, then, leads to an earnings adjustment in response to the new marginal tax rate. The post-tax price of earnings (or retention rate) $1-\tau_i$ increases by $\mathcal{T}'(y_i^*) - \mathcal{T}'(y_i^* - \partial_L y_i^{\text{int}}) = -\mathcal{T}''(y_i^*)\partial_L y_i^{\text{int}}$. This change in the marginal price of earnings applies only to the last, marginal dollar. Therefore, the labor supply response to the marginal price change is governed by the *compensated* substitution effect $\frac{\partial y^*}{\partial(1-\tau)} |_{u=\overline{u}}$.[2] Putting these two results together and taking the average over all households with earnings y implies that the intensive-margin effect is proportional to $\mathbb{E}_y \frac{\eta_i}{1+\theta_i \zeta_i^c}$.

One useful special case is when the tax function $\mathcal{T}(\cdot)$ takes the form $\mathcal{T}(y) = y - \text{const.} \times y^{1-\theta}$, where θ is a parameter capturing the progressivity in the tax code. It is immediate to verify that if the tax function takes this form then $\theta_i = \theta$ for all i. This tax function is commonly used in public finance. Heathcote et al. (2017) showed that it is a good approximation of tax rates in the U.S. and estimated $\theta = 0.18$. The best estimate for Norway (2015) is $\theta = 0.15$.

Changes in marginal tax rates

In a model with a linear income tax and no extensive margin responses to tax changes there would be a one to one relationship between the observed earnings response and the uncompensated earnings elasticity in the Slutsky equation. In a model with heterogeneity, non-linear taxes and extensive margin decisions, the connection between earnings responses to tax changes and the uncompensated earnings elasticity is more involved. In order to use earnings responses to make inference about the Slutsky parameters one must factor out extensive margin responses. Another complication is that with a non-linear income tax, a small increase in the marginal tax rate at income level y implies an increase in the average tax rate of those who earn income above y. The earnings response to such a change in the tax rate will be a combination of behavioral responses to changes in the marginal and average tax rates.

Our focus from now on is a piece-wise linear income tax scheme since the Norwegian and most other real-world income tax schemes take this form. In particular, we characterize $\mathcal{T}(\cdot)$ as a collection $\{\tau_j, \underline{y}_j, \underline{\mathcal{T}}_j\}_{j=1}^J$, where τ_j is the tax rate in bracket j, $\underline{y}_j$ is the income level at which this tax rate kicks in, and $\underline{\mathcal{T}}_j$ is the tax an individual with income $\underline{y}_j$ has to pay.[3] We can then express the taxes paid by individuals with earnings y that fall into bracket j as

$$\mathcal{T}(y) = \tau_j(y - \underline{y}_j) + \underline{\mathcal{T}}_j. \tag{2.11}$$

We now turn to analyzing the earnings responses to changes in the marginal tax rates across different brackets. The simplest case to analyze is when the marginal tax rate increases at the top income bracket.

[2] See also Saez (2001) for a simple graphical proof of this fact.

[3] The triplets that characterize the tax brackets are related to each other in a recursive way: $\underline{\mathcal{T}}_{j+1} = \underline{\mathcal{T}}_j + \tau_j(\underline{y}_{j+1} - \underline{y}_j)$.

A change in the marginal tax rate at the top bracket

A small increase in the top marginal tax rate may have both intensive and extensive margin effects on earnings. The intensive margin effect implies an adjustment of earnings within the top bracket. This adjustment depends on how earnings respond to changes in unearned income and to changes in the relative price of consumption (the net of tax rate). To see this, note that a small increase in the marginal tax rate applied to income above $\underline{y}_j$ creates a substitution effect and an income effect that depends on how far above $\underline{y}_j$ a person earns. For someone who earns exactly $\underline{y}_j$ prior to the increase in the top rate there is no income effect, only a substitution effect. In contrast, someone who earns far above $\underline{y}_j$ will suffer a considerable income loss if the bracket tax rate increases. This person's intensive margin earnings response will therefore be a combination of the substitution and income effect. In the event of an extensive margin response, this person's earnings are reduced to zero. The extensive margin effect is therefore proportional to the marginal propensity to stop working.

Taken together, the change in average earnings in the top bracket due to an increase in the top marginal tax rate is given by

$$\partial_{\tau_J} \mathbb{E}_{\geq \underline{y}_J} y_i^* = \mathbb{E}_{\geq \underline{y}_J} \left[(\underline{y}_J - y_i^*)\eta_i - \frac{y_i^*}{(1-\tau_J)} \zeta_i^c \right] - \int_{\underline{y}_J}^{\infty} y \lambda_{\tau_J}(y) \frac{dH(y)}{1 - H(\underline{y}_J)}, \tag{2.12}$$

where $\lambda_{\tau_J}(y) \geq 0$ is the marginal propensity to stop working in response to an increase in the top rate τ_J.

Changing the marginal tax rate at lower income brackets

An increase in the marginal tax rate at some lower bracket $j < J$ will have an effect on earnings decisions for those in bracket j, but it also affects the average tax rate for those with earnings above bracket j. The response to such a tax increase may again be an interior adjustment of earnings, or it may induce marginal workers to leave the labor market. The earnings response within bracket j is similar to the response within the top-bracket when the top rate increases. The intensive-margin response is driven both by a substitution and an income effect. The change in average earnings for individuals with earnings in bracket $j < J$ is given by

$$\partial_{\tau_j} \mathbb{E}_{\in \left[\underline{y}_j, \underline{y}_{j+1}\right]} y_i^* = \mathbb{E}_{\in \left[\underline{y}_j, \underline{y}_{j+1}\right]} \left[\left(\underline{y}_j - y_i^* \right) \eta_i - \frac{y_i^*}{(1-\tau_j)} \zeta_i^c \right] - \int_{\underline{y}_j}^{\underline{y}_{j+1}} y \lambda_{\tau_j}(y) \frac{dH(y)}{H(\underline{y}_{j+1}) - H(\underline{y}_j)}. \tag{2.13}$$

For those who earn above bracket j, however, there is no price change and no substitution effect since the marginal tax rate in their bracket is unchanged. The interior response here is driven only by a negative income effect. The change in average earnings conditional on earning above $\underline{y}_{j+1}$ due to a unit increase in the marginal tax rate in bracket j is therefore given by

$$\partial_{\tau_j} \mathbb{E}_{\geq \underline{y}_{j+1}} y_i^* = \mathbb{E}_{\geq \underline{y}_{j+1}} \left[(\underline{y}_j - \underline{y}_{j+1})\eta_i \right] - \int_{\underline{y}_{j+1}}^{\infty} y \lambda_{\tau_j}(y) \frac{dH(y)}{1 - H(\underline{y}_{j+1})} \tag{2.14}$$

2.3 The deadweight loss of income taxation

Tax induced changes in behavior have an impact on economic efficiency. An income (or consumption) tax drives a wedge between what an hour of work is worth to an employer, and the income the employees can take home. This price distortion may prevent mutually advantageous trade from taking place and if it does, it creates a deadweight loss (DWL). The DWL associated with a tax (or a tax reform) is the welfare loss, measured in money, that exceeds the revenue collected by the government.

There are alternative ways to measure this loss. Our analysis applies the measure hinted at above, which is based on a compensating variation logic: *The deadweight loss of an income tax is the difference between the income compensation a worker requires in order to be as well off as he/she was without the tax and the revenue collected by the tax.*[4]

Above we saw that an increase in the income tax rate may generate behavioral responses both at the intensive and extensive margin. Both types of responses will, potentially, create deadweight losses. This section focuses on how an increase in the tax rate affects the deadweight loss at the intensive margin, but includes a brief discussion of how extensive margin responses affect the calculation of deadweight losses.

The deadweight loss of a linear income tax

We begin by considering the simplest case with a linear income tax $\mathcal{T}(y) = \tau y$. The compensating variation (CV) is the lump sum payment needed to make a person in a world with taxes equally well off as if there were no taxes. One way to quantify the compensating variation of a marginal (and average) tax rate τ is by using the indirect utility function,

$$v_i(\tau, R_i) \equiv \max_{c_i, y_i}\{u_i(c_i, y_i) \mid c_i \leq y_i(1-\tau) + R_i\}. \tag{2.15}$$

The compensating variation for individual i associated with a tax rate τ is then implicitly defined by the equation $v_i(\tau, R_i + \text{CV}) - v_i(0, R_i) = 0$.

An alternative way to get at the compensating variation is to use the expenditure function. For an individual i the expenditure function gives the minimum expenditure needed in order to attain utility $\overline{u}$ when income is taxed at rate τ and unearned income is R_i,

$$E_i(\tau, \overline{u}) \equiv \min_{c_i, y_i}\{c_i - (1-\tau)y_i - R_i \mid u_i(c_i, y_i) \geq \bar{u}\}.$$

The solution to this minimization problem defines the compensated, or Hicksian, demand and earnings functions $c_i^c(\tau, \overline{u})$ and $y_i^c(\tau, \overline{u})$.[5] The compensating variation is then the difference between the expenditure

[4] Equivalent variation is another measure of the deadweight loss associated with taxation. This measure asks how much income a worker would be willing to give up in order to get rid of the tax. See, e.g., Auerbach and Hines Jr (2002) for a general discussion of efficiency and taxation.

[5] Since we discuss economic efficiency in a static framework we return to the notation introduced in the textbook model above, but allow for heterogeneity in the utility function and unearned income.

needed to attain the same utility in a world with and without income tax,

$$\text{CV}_i(\tau, \overline{u}) \equiv E_i(\tau, \overline{u}) - E_i(0, \overline{u}). \tag{2.16}$$

To determine the deadweight loss, we then subtract from the compensating variation the tax revenue that the government receives if taxpayers were fully compensated. With a linear, proportional income tax, the compensated tax revenue is given by $\tau y_i^c(\tau, \overline{u})$. Thus, the deadweight loss associated with introducing a proportional income tax τ is given by

$$\text{DWL}_i(\tau, \overline{u}) \equiv E_i(\tau, \overline{u}) - E_i(0, \overline{u}) - \tau y_i^c(\tau, \overline{u}). \tag{2.17}$$

Expression (2.17) measures the efficiency loss that occurs when a linear income tax is introduced. A related, and arguably more relevant question is how the deadweight loss changes when the government alters an already existing tax rate. By differentiating expression (2.17) with respect to the tax rate, we obtain the marginal deadweight loss (MDWL) of individual i,

$$\text{MDWL}_i \equiv \frac{\partial E_i(\tau, \bar{u})}{\partial \tau} - y_i^c(\tau, \bar{u}) - \tau \frac{\partial y_i^c(\tau, \bar{u})}{\partial \tau} = -\tau \frac{\partial y_i^c(\tau, \bar{u})}{\partial \tau} = \frac{\tau}{1-\tau} y_i^* \zeta_i^c. \tag{2.18}$$

Expression (2.18) shows that the larger the compensated elasticity, the larger the marginal deadweight loss from income taxation. It highlights the need to have an estimate of the substitution effect in order to assess the efficiency costs of income taxation.

It is useful to normalize the marginal deadweight loss by the additional tax revenue that is generated from a small increase in the tax rate. The marginal tax revenue (MTR) collected from individual i is given by

$$\text{MTR}_i = y_i^* + \tau \frac{\partial y_i^*}{\partial \tau} = y_i^* - \frac{\tau y_i^*}{1-\tau} \zeta_i^u, \tag{2.19}$$

where the first term captures the mechanical increase in taxes, while the second term accounts for changes in the tax collected due to the individual's behavioral response. Dividing (2.18) by (2.19) and substituting the uncompensated elasticity using the Slutsky equation (2.3), we then obtain the marginal deadweight loss per unit of tax revenue,

$$\frac{\text{MDWL}_i}{\text{MTR}_i} = \frac{\frac{\tau}{1-\tau} y_i^* \zeta_i^c}{y_i^* - \frac{\tau y_i^*}{1-\tau} \left(\zeta_i^c + (1-\tau)\eta\right)}. \tag{2.20}$$

Expression (2.20) shows that the normalized marginal deadweight loss increases in the compensated elasticity and declines in the absolute value of the income effect η. A larger compensated elasticity implies that earnings are sensitive to the net of tax rate, and consequentially the economic surplus that is lost is large. A high absolute value of the income effect implies that workers are sensitive to the income they lose with increased taxation and will cover a substantial part of this loss by increasing their earnings. While this income effect leaves the marginal deadweight loss unchanged, it increases the tax revenue that the government collects and lowers the deadweight loss per unit of tax revenue.

Efficiency losses from increasing marginal tax rates

In the previous paragraph we developed an expression for the additional efficiency loss generated by a change in the linear tax rate. However, most governments use a progressive, piece-wise linear tax scheme to tax labor earnings. Applying a similar logic as above, we can find an expression for the normalized marginal deadweight loss from an increase in the tax rate in bracket j. However, there are two key differences worth noting.

First, an increase in τ_j creates a substitution effect only for those who are located in bracket j since the change in the marginal price of earnings applies only to the last, marginal dollar. Second, a higher tax rate does not affect the income for anyone in brackets below j, but generates a loss in income for anyone with earnings above $\underline{y}_j$. The reduction in income of a person in bracket j depends on how far above $\underline{y}_j$ she earns, and amounts to $(y_i^* - \underline{y}_j)d\tau_j$. By comparison, anyone located in a bracket above j experiences the same mechanical drop in income equal to $(\underline{y}_{j+1} - \underline{y}_j)d\tau_j$.

These considerations have three implications for the calculation of the efficiency loss associated with this type of reform. First, and most obvious, anyone located in brackets below j is not affected by the reform, and consequently, does not contribute to any further loss in efficiency,

$$\text{MDWL}_i = 0 \quad \text{for all } i \text{ who earn } y_i^* < \underline{y}_j. \tag{2.21}$$

Second, the reform creates an income effect (but no substitution effect) for every person located in a bracket above j. The loss in income matches exactly the compensation these individuals must receive in order to attain their pre-reform utility. Therefore, their *compensated* labor supply does not respond to an increase in the tax rate. In other words, the tax change is nondistortinary for individuals in bracket above j, and therefore,

$$\text{MDWL}_i = 0 \quad \text{for all } i \text{ who earn } y_i^* > \underline{y}_{j+1}. \tag{2.22}$$

Third, a higher marginal tax rate in bracket j creates additional deadweight loss for anyone in bracket j due to the substitution effect. Analogous to the steps taken in deriving expression (2.20), we obtain the following formula for the marginal deadweight loss per unit of additional tax revenue,

$$\frac{\text{MDWL}_i}{\text{MTR}_i} = \frac{\frac{\tau_j}{1-\tau_j} y_i^* \zeta_i^c}{y_i^* - \underline{y}_j + \tau_j \left(\eta_i (\underline{y}_j - y_i^*) - \frac{y_i^*}{1-\tau_j} \zeta_i^c \right)} \quad \text{for all } i \text{ who earn } y_i^* \in \left[\underline{y}_j, \underline{y}_{j+1} \right]. \tag{2.23}$$

Efficiency losses at the extensive margin

An increase in the marginal tax rate may induce some workers to leave the labor market. Exit or entry decisions are based on a comparison of utility when participating (and choosing earnings optimally) and the utility one obtains if not participating. A person who leaves the labor market in response to a small increase in the tax rate was initially indifferent between participating or not, and, consequentially, needs no compensation to attain the pre-reform utility.

However, the decision to leave the labor market may lead to a sharp drop in government net revenues. When a person with income y_i^* drops out of the labor market, the reduction in government revenues amounts to $\mathcal{T}(y_i^*) - \mathcal{T}(0)$, where $\mathcal{T}(0) \leq 0$ are government transfers to individuals with zero earnings. Thus, there is a deadweight loss of $\mathcal{T}(y_i^*) - \mathcal{T}(0)$ for each person with income $y_i^* \geq \underline{y}_j$ that leaves the labor market because of an increase in the tax rate in bracket j. The change in the expected deadweight loss due to extensive margin responses following a small increase in the marginal tax rate τ_j is therefore given by

$$\partial_{\tau_j} \mathbb{E}_{\geq \underline{y}_j} \mathrm{DWL}_i^{\mathrm{ext}} = \int_{\underline{y}_j}^{\infty} \left(\mathcal{T}(y) - \mathcal{T}(0)\right) \lambda_{\tau_j}(y) \frac{dH(y)}{1 - H(\underline{y}_j)}, \tag{2.24}$$

where $\lambda_{\tau_j}(y)$ is the marginal propensity to stop working in response to an increase in the tax rate τ_j.

2.4 Optimal tax rates

The deadweight loss of an increase in the marginal tax rate is one of the arguments that matter for the optimal tax rates. Another argument that matters for the calculation of optimal tax rates is the distribution consequences. In an economy with income inequality and progressive taxation, optimal tax rates depend on the welfare weights policy makers assign to individuals with different income levels. Thus, any statement about optimal taxation involves a trade-off between the concerns about efficiency and redistribution.

2.4.1 Optimal tax rates at the top

Suppose that the government assigns a welfare weight $g_J < 1$ to individuals who belong to the top income bracket. This means that the government is indifferent between g_J more dollars of tax revenue and one more dollar consumed by earners in the top bracket. An optimizing government will then choose the top marginal tax rate that maximizes $\mathcal{R}(\tau_J) - g_J \mathcal{M}(\tau_J)$, where $\mathcal{R}$ is tax revenue collected with a rate τ_J and $\mathcal{M}$ is a monetary value of the welfare loss experienced by those in the top bracket if they are taxed at rate τ_J. The first order condition for this problem is $\frac{\partial \mathcal{R}(\tau_J)}{\partial \tau_J} - g_J \frac{\partial \mathcal{M}(\tau_J)}{\partial \tau_J} = 0$.

The expression for the marginal tax revenue in the top income bracket is given by the denominator in equation (2.23). An increase in the marginal tax rate gives a mechanical increase in tax income in addition to an effect that depends on how top income earners adjust earnings in response to a tax increase.[6] The monetary value of the welfare loss experienced by taxpayers in the top bracket is exactly matched by the mechanical tax revenue effect for the government. The fact that taxpayers adjust their earnings behavior in response to a change in the marginal tax rate does not have a first-order effect on their utility if we assume - as we do - they had optimized earnings initially (the envelope condition). Based on this, and assuming homogeneous response elasticities among earners in the top income bracket, the optimal top income tax rate is such that (see Saez (2001) for a formal derivation)

$$\frac{\tau_J}{1 - \tau_J} = \frac{(1 - g_J)(a - 1)}{a\zeta_J^c + (1 - \tau)\eta_J (a - 1)}, \text{ with } a \equiv \mathbb{E}_{\geq \underline{y}_J} \left[\frac{y_i}{\underline{y}_J} \right]. \tag{2.25}$$

The optimal top income tax rate (2.25) is decreasing in the compensated elasticity, and increasing in the

[6] As in Saez (2001), we focus on earnings adjustments along the intensive margin and ignore the possibility that a high income earner stops working in response to an increase in the top tax rate.

absolute value of the income effect. To understand why this is the case, recall that a larger substitution effect implies larger efficiency costs from taxation. Therefore, the tax rate must be lower in optimum. In contrast, a larger absolute value of the income effect leads to a higher optimal tax rate since raising the tax rate reduces disposable income and hence increases labor supply and tax revenue.

It is important to realize that the LHS of equation (2.25) does not provide an explicit closed form expression of the optimal tax rate. We can see that directly from the observation that the tax rate appears also on the RHS of this optimality condition. In addition, and perhaps more fundamentally, neither the relevant behavioral responses, nor the average income in the top bracket, are entities that are tax invariant. On the contrary, it is likely that they depend critically on the marginal tax rate.

2.4.2 Optimal tax rates at lower brackets

The previous paragraph considered only the problem of optimal tax rates at the top end of the income distribution. In this section, we investigate the issue of optimal tax rates in any given tax bracket $j \leq J$. As in previous research, we will ignore the possibility that an increase in the marginal tax rate at lower brackets creates an income effect that is large enough to push a person into a higher tax bracket.

Suppose there exists a fixed set of tax brackets, indexed by $j = 1, 2, ..., J$, and let $\boldsymbol{\tau} = (\tau_1, \tau_2, ..., \tau_J)$ denote the corresponding tax rates. The government assigns a welfare weight $g(y)$ to individuals with income y and solves

$$\max_{\boldsymbol{\tau}} \int \left(\mathcal{R}(\boldsymbol{\tau}, y) - g(y)\mathcal{M}(\boldsymbol{\tau}, y)\right) dH(y). \tag{2.26}$$

The function $\mathcal{R}(\boldsymbol{\tau}, y)$ captures the tax revenue collected from individuals with income y at tax rates $\boldsymbol{\tau}$, and $\mathcal{M}(\boldsymbol{\tau}, y)$ is the monetary value of the welfare loss experienced by individuals with income y if they are taxed with rates $\boldsymbol{\tau}$. The first-order conditions for the tax rates are given by

$$\int \left(\frac{\partial \mathcal{R}(\boldsymbol{\tau}, y)}{\partial \tau_j} - g(y)\frac{\partial \mathcal{M}(\boldsymbol{\tau}, y)}{\partial \tau_j}\right) dH(y) = 0 \quad \text{for all } j = 1, 2, ..., J. \tag{2.27}$$

Just as it was the case for top income tax rate, expression (2.27) shows that an optimal tax rate in bracket j must necessarily be set such that the marginal tax revenue equals the marginal social welfare loss.

Additional tax revenue is generated both through a mechanical and a behavioral effect. Under the assumption that response elasticities are homogeneous across individuals with the same income, we can express the average marginal tax revenue as

$$\int \frac{\partial \mathcal{R}(\boldsymbol{\tau}, y)}{\partial \tau_j} dH(y) = \int_{\underline{y}_j}^{\underline{y}_{j+1}} \left[\underbrace{(y - \underline{y}_j)}_{(a)} + \underbrace{\frac{\tau_j}{1-\tau_j}\Big((1-\tau_j)\eta(y)(\underline{y}_j - y) - y\zeta^c(y)\Big)}_{(b)} \right] dH(y) \tag{2.28}$$

$$+ \int_{\underline{y}_{j+1}}^{\infty} \left[\underbrace{(\underline{y}_{j+1} - \underline{y}_j)}_{(c)} + \underbrace{\tau(y)\eta(y)(\underline{y}_j - \underline{y}_{j+1})}_{(d)} \right] dH(y),$$

where $\tau(y)$ denotes the marginal tax rate at income level y. To understand why expression (2.28) takes this form, consider a perturbation of the tax rate τ_j. This changes the marginal and average tax rate for individuals located in bracket j, and the average tax rate for everyone earning more than $\underline{y}_j$ - including those in tax brackets above j. This implies that a tax rate change induces two different kinds of earnings adjustments depending on whether a person is located in bracket j or above. For someone located in tax bracket j, a change in τ_j creates a substitution effect and an income effect that depends on how far above $\underline{y}_j$ a person earns. In contrast, this tax rate change creates only an income effect for someone in a higher tax bracket.[7] These behavioral effects on government revenue are reflected in terms (b) and (d) in expression (2.28).

In addition to the behavioral effects, there are also pure mechanical effects on tax receipt. For someone located in bracket j, the mechanical effect depends on his earnings relative to $\underline{y}_j$, whereas for someone located in a higher tax bracket, the mechanical effect is constant and equal to the width of tax bracket j. These two mechanical effects are captured by terms (a) and (c), respectively.

The monetary value of the welfare loss experienced by taxpayers in brackets j and above is exactly matched by the mechanical tax revenue effect for the government. The average marginal social welfare loss is therefore given by

$$\int g(y)\frac{\partial \mathcal{M}(\boldsymbol{\tau},y)}{\partial \tau_j}dH(y) = \int_{\underline{y}_j}^{\underline{y}_{j+1}} g(y)(y-\underline{y}_j)dH(y) + \int_{\underline{y}_{j+1}}^{\infty} (\underline{y}_{j+1}-\underline{y}_j)g(y)dH(y). \tag{2.29}$$

By putting the two terms (2.28) and (2.29) together, we find that an optimal tax rate τ_j in bracket $j = 1, 2, ..., J$ must satisfy

$$\frac{\tau_j}{1-\tau_j} = \frac{\mathbb{E}_{\in\left[\underline{y}_j,\underline{y}_{j+1}\right]}\left[(1-g(y))(y-\underline{y}_j)\right] + \Lambda_j(\underline{y}_{j+1}-\underline{y}_j)\mathbb{E}_{\geq \underline{y}_{j+1}}\left[(1-g(y))+\tau(y)\eta(y)\right]}{\mathbb{E}_{\in\left[\underline{y}_j,\underline{y}_{j+1}\right]}\left[y\zeta^c(y)+(1-\tau_j)\eta(y)(y-\underline{y}_j)\right]}. \tag{2.30}$$

The factor $\Lambda_j \equiv \left(1-H(\underline{y}_{j+1})\right) / \left(H(\underline{y}_{j+1})-H(\underline{y}_j)\right)$ captures the mass of individuals in higher brackets relative to the mass of individuals in bracket j. In the case of the top bracket, $\Lambda_J = 0$ and it is easy to verify that expression (2.30) simplifies to expression (2.25) derived above. For tax brackets below the top, the additional term in the numerator captures the mechanical (net of the welfare loss) and behavioral effects on tax revenue in brackets above.

2.5 Key challenges taking the theory to the data

The theory above defines the parameters of interest for empirical analyses of labor supply and income taxation. A key problem remains: taking the theory to the data.

Based on the labor supply model outlined above, the early econometric literature on labor supply speci-

[7] We ignore the possibility that an increase in the marginal tax rate at lower brackets creates an income effect that is large enough to push a person into a higher tax bracket. Furthermore, we restrict ourselves to earnings adjustments along the intensive margin.

fied a regression equation of the type

$$S_i = \beta_0 + \beta_1\omega_i + \beta_2 R_i + \beta_3 x_i + \nu_i,$$

where S is a measure of labor supply (typically hours of work), ω is net-of-tax wage and R unearned income. With exogenous variation in ω and R, β_1 captures how a wage increase affects labor supply holding unearned income fixed, and β_2 how an increase in unearned income affects labor supply holding wages fixed. Given consistent estimates of these parameters one could use the Slutsky equation to try to recover the compensated effect on labor supply.

A key challenge is that with observational data it is difficult to obtain exogenous variation in net wages or unearned income. Observed variation in wages, unearned income and hours worked will depend on other unobserved factors that matter for labor supply (and demand). Wages and unearned income will therefore be correlated with the error term ν_i.

Another problem is that we do not observe wages for those who choose not to participate in the labor market. Selective participation in the labor market creates an endogeneity problem that makes it difficult to obtain consistent estimates of the intensive margin labor supply. The solution to this problem is either to make strong parametric assumptions about wage offers and participation costs or, if available, use instruments that influence participation but not wages.

In fact, heterogeneity in the "taste for work" can bias the intensive margin elasticities even without the non-participation problem. The reason is that persons with a low cost of effort or high taste for work will typically have both higher wages and put in more work hours than those with a lower taste of work. In this case, regressing hours worked on wages would overestimate how elastic individuals are to changes in wages or taxes.[8]

To circumvent endogeneity problems in labor supply regressions, researchers have turned to arguably exogenous variation in wages net of taxes coming from changes in the marginal tax rate. In practise, income tax schemes are typically progressive and piece wise linear. Such a tax function creates variation in marginal tax rates and therefore in wages net of taxes. There are, however, several problems associated with using this variation directly to estimate the labor supply elasticity. One obvious issue is that it creates a reverse causality problem, since the marginal tax rate is an increasing function of labor income. As discussed in the next section, a possible solution to this problem is to use tax reforms as instruments.

Two other more subtle problems also arise if the the income tax is non-linear, but the econometric model assumes that each individual faces a constant marginal tax rate. One issue is that a non-linear tax creates convex kinks in the budget set at the income levels where the marginal tax rate increases. The presence of kinks may lead one to underestimate the labor supply elasticity in a model that assumes constant marginal taxes. Another issue is that a linear model mis-specifies the income of a person in higher income tax brackets. Hausman's solution to that problem was to introduce a virtual income equal to the difference between the

[8] A considerable body of work assumes that the individuals' unobserved taste for work is uncorrelated with productivity. If one in addition makes a parametric assumption about the distribution of the taste for work (e.g. that it is extreme value distributed), it is possible to recover individuals preferences from observed variation in wages. This approach forms the basis for the microsimulation studies of labor supply in Norway.

person's tax liability in a linear system in which only the highest marginal tax rate applies and the person's actual tax liability.

In the empirical analyses below, we will deal with the above challenges by obtaining plausibly exogenous variation in wages net of taxes and unearned income, by distinguishing between extensive and intensive margin responses, and by accounting for non-linearities in the tax system both in the estimation of the earnings responses and in how these responses are interpreted.

3 Uncompensated responses to tax reforms

3.1 Identification of uncompensated responses using tax reforms

Since Feldstein (1995), a large body of empirical work has attempted to quantify behavioral responses to (changes in) marginal tax rates by estimating the elasticity of taxable income (ETI). The starting point is a log-linear regression specification of the following form;

$$\log(y_{it}) = \alpha + \text{ETI} \cdot \log(1 - \tau_{it}) + \nu_t + \mu_i + \varepsilon_{it}, \tag{3.1}$$

where y_{it} is reported taxable income of individual i in year t, τ_{it} is the marginal tax rate, ν_t denotes time effects, μ_i is an individual fixed effect, and ε_{it} is an error term. The typical estimating equation is some version of equation (3.1) in difference form;

$$\Delta\log(y_{it}) = \text{ETI} \cdot \Delta\log(1 - \tau_{it}) + \Delta\nu_t + \Delta\varepsilon_{it}, \tag{3.2}$$

where Δ refers to differences between the periods t and $t - s$, with $s > 0$.

Estimating the ETI using a specification like equation (3.2) is complicated due to two main empirical challenges. First, in a progressive tax system, the marginal tax rate τ_{it} and taxable income y_{it} are jointly determined, creating a correlation between the net-of-tax rate change $\Delta\log(1-\tau_{it})$ and the error term $\Delta\varepsilon_{it}$. To address this challenge, most studies use an instrument for $\Delta\log(1 - \tau_{it})$ that is based on "mechanical" changes in tax rates driven by changes in tax schedules. Following Weber (2014) and Vattø (2020), we define the mechanical change in the net-of-tax rate as follows;

$$\Delta\log(1 - \tau_{it})^{mech} = \log(1 - \tau_{it}(y^*_{it-s-1})) - \log(1 - \tau_{it-s}(y^*_{it-s-1})), \tag{3.3}$$

where y^*_{it-s-1} in the first and second term on the right-hand side denotes lagged base-year income inflated to years t and $t - 3$, respectively, using growth in median income. By using the mechanical change in the net-of-tax rate as an instrument for the observed change in the net-of-tax rate, one isolates changes in tax rates due to changes in the tax schedule.

The second main empirical challenge in the ETI literature is related to a potential correlation between (lagged) base-year income and the error term, which might occur for at least two reasons: (i) different income growth rates across the population (e.g. larger income growth for high-income earners), and (ii) reversion to the mean. The existing literature has attempted to address these challenges by controlling for demographic characteristics and some function $f(\cdot)$ of (lagged) base-year income, and by imposing sample restrictions to exclude low-income earners, since mean-reversion tends to be most pronounced at the bottom of the income distribution.

Our baseline specification follows common practice in the literature: We use a difference length of $s = 3$, and account for differential income growth by including a third-order polynomial in lagged base-year income. We also report a number of robustness checks to show that our main results are robust to including different specifications of $f(\cdot)$, controlling for demographic characteristics, and excluding low-income earners.

3.2 Data, sample, and variation in tax rates

3.2.1 Data and sample

Our empirical analyses are based on several administrative data sources that we are able to link together using unique identifiers for individuals and households. This results in a matched panel dataset covering the full Norwegian population in the period 1995–2018. The dataset includes detailed information from income tax returns as well as individual characteristics such as age, sex, educational attainment, marital status, and number of children.

We construct our baseline sample using wage earners between 25 and 61 years old, with wage earners being defined as individuals for whom wage income is the main source of labor income. We exclude students and individuals receiving pensions or unemployment benefits. With these restrictions, our baseline estimation sample contains about 18.7 million observations and 2.6 million individual wage earners. Summary statistics for our baseline estimation sample are provided in Table 3.1.

Marginal tax rates are not directly observed in tax return data. We therefore simulate marginal tax rates for each taxpayer based on information from tax returns and a detailed tax simulation model of the Norwegian income tax system. Based on an extended version of the model used in Vattø (2020), we compute the marginal tax rate τ_{it} for each individual in the sample by increasing taxable income by a small amount (NOK 500). In particular, we compute the marginal tax rate for an individual with income $\tilde{y}$ as $\tau(\tilde{y}) = [T(\tilde{y} + 500) - T(\tilde{y})]/500$, where $T(\cdot)$ is tax liability.

3.2.2 The Norwegian tax system and tax reforms

Norway has a dual income tax system with a flat tax on capital income and a progressive tax on labor income. During our sample period from 1995 to 2018, labor income was taxed at a flat rate ranging between 23 and 28 percent in addition to a social security contribution rate of 7.8 percent (8.2 percent from 2014), and according to a surtax rate schedule with between one and four different tax rate tiers.[9]

We distinguish between three different periods with changes in the surtax schedule affecting different groups of workers: (i) the 2000 reforms, with the largest changes in tax rates for high-income earners; (ii) the 2006 reform, with changes primarily around the middle of the income distribution; and (iii) the 2016

[9] Whereas the tax base for paying the flat tax rate is subject to a number of deductions, the base for paying surtax rates and social security contributions consists of total labor income without any deduction possibilities.

Table 3.1: Summary statistics for the baseline estimation sample (1999–2018)

Variable	Mean	Std. Dev.	p25	p50	p75
		Demographics			
Age	41.93	8.95	34	42	49
Number of children	0.81	1.02	0	0	2
Years of education	12.44	2.83	10	12	15
Male (percent)	55.6				
Married (percent)	55.6				
		Income and taxes			
Taxable income	507,983	316,760	351,194	450,964	590,089
Income taxes	149,901	198,739	83,418	116,502	171,124
Marginal tax rate (percent)	40.2	7.7	35.8	35.8	47.8
		Key variables in the regressions			
$\Delta\log(y_{it})$	0.139	0.345	0.047	0.129	0.230
$\Delta\log(1-\tau_{it})$	0.008	0.098	0	0.004	0.011
$\Delta\log(1-\tau_{it})^{mech}$	0.015	0.057	0	0.004	0.011
Observations	18,707,745				
Individuals	2,600,580				

Notes: This table reports summary statistics for the sample used in the estimation of our baseline specification. Δ refers to differences between the periods t and $t-3$, and level variables refer to the base-year, year $t-3$. Monetary values are reported in 2018 NOK.

reform, with the largest changes in the lower parts of the distribution. During the period from 2010 to 2013, the surtax schedule remained unchanged, except for annual adjustments of tax bracket cut-offs in accordance with expected average wage growth. We exclude these years from our estimation sample, as there is no first stage in this period.

The main changes in the surtax schedule are illustrated in Figure 3.1. For each of the three reform periods, we contrast the post-reform tax schedule with the tax schedule from three years before, in line with the three-year differences in income and tax rates used in our baseline empirical specification.

Figure 3.1: Changes in tax schedules for three different reform periods

(a) 2000 reform (b) 2006 reform (c) 2016 reform

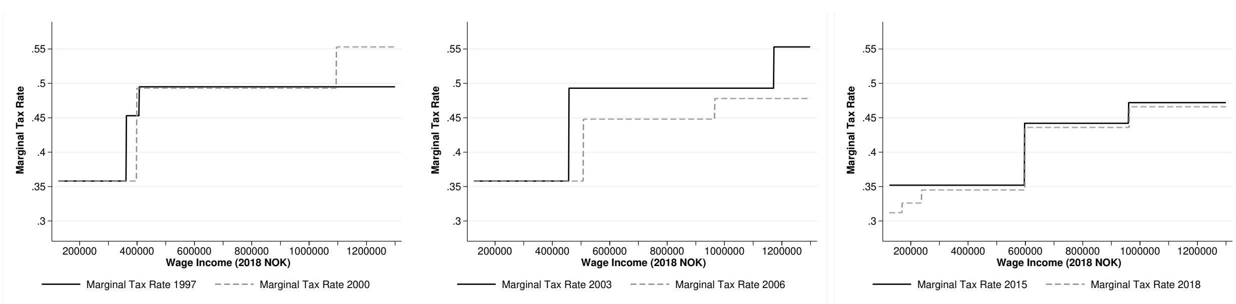

Notes: This figure shows the marginal tax rates that apply to different income tax brackets, for three different reform periods. The solid black lines show pre-reform tax schedules, while the dashed grey lines show post-reform schedules.

The 2000 reform changed the two-tier surtax schedule in two ways, with substantial changes in both the surtax rates and in bracket cut-offs. The surtax rates were increased from 9.5 and 13.7 to 13.5 and 19.5 percent for the first and second bracket, respectively. At the same time, the bracket cut-off for the second bracket was considerably increased: while the top tax rate applied to more than 30 percent of all wage earners in 1997, it was only levied on the top 1-2 percent of the income distribution in 2000 (see Vattø, 2020). For these high-income earners, the marginal tax rate increased from 49.5 in 1997 to 55.3 in 2000.

In the 2006 reform, the surtax rates were reduced from 13.5 and 19.5 to 9 and 12 percent, respectively. The bracket cut-off for the second bracket was lowered, but with the large reductions in surtax rates, a majority of workers in the upper half of the income distribution nevertheless experienced substantial reductions in the marginal tax rate as a result of this reform.

The 2016 reform changed the surtax rate schedule from a two-tier to a four-tier schedule, with the third and fourth tiers replacing the first and second tiers in the pre-reform system. This reform was introduced in the middle of a period with reductions in the flat tax rate; from 28 percent in 2013 to 22 percent in 2019. In sum, the changes in the surtax schedule and the changes in the flat tax rate implied small reductions in marginal tax rates for most workers, with the largest changes in the lower part of the income distribution.

3.3 Empirical results

This section presents estimates of the ETI in terms of results from estimating equation (3.2) using 2SLS, instrumenting the change in the net-of-tax rate with the mechanical change in the net-of-tax rate defined in equation (3.3). We begin by presenting results from our baseline specification estimated on a stacked panel dataset including all the relevant three-year differences in the period 1999–2018. We then present a number of robustness checks showing that our main results are robust to different ways of controlling for differential income growth, to excluding low-income earners, to different definitions of the instrument, and to varying the length of the intervals used to measure changes in taxable income and net-of-tax rates. We also present results from our baseline specification estimated on stacked panel datasets covering the different reform periods discussed in Section 3.2.2.

3.3.1 Baseline results

The first column of Table 3.2 shows results from our baseline specification, which accounts for differential income growth by including a third-order polynomial in lagged base-year income. We estimate equation (3.2) using 2SLS, instrumenting the change in the net-of-tax rate with the mechanical change in the net-of-tax rate defined in equation (3.3). The regression is weighted by lagged base-year income.

The table shows that the first stage of the baseline regression is very strong, with a significant first-stage coefficient and a large first-stage F-statistic, indicating that weak instruments is not a concern in this setting. The estimated baseline ETI of 0.13 is also statistically significant, and implies a 1.3 percent increase in taxable income in response to a 10 percent increase in the net-of-tax rate.

Our baseline ETI estimate is broadly comparable to estimates in the existing literature. Saez et al. (2012) review the literature and conclude that estimates of the ETI from the U.S. (after Feldstein, 1995) range from 0.12 to 0.4, while Neisser (2021) finds that average ETIs (for income before deductions) are in the range of 0.05 to 0.12, based on a meta-regression analysis covering 61 studies from both the U.S. and several European countries. Our results are also in line with recent studies for wage earners in Norway: Thoresen and Vattø (2015) and Vattø (2020) report ETIs of about 0.05 and between 0.11 and 0.15, respectively.

Columns (2)–(7) of Table 3.2 show results from estimating equation (3.2) with different approaches to accounting for differential income growth. We explore a similar set of alternatives as in Kleven and Schultz (2014): no income controls; controlling for the log of lagged base-year income; controlling for a third-order polynomial in the log of lagged base-year income (our baseline specification); controlling for a ten-piece spline in the log of lagged base-year income; controlling for a third-order polynomial in the log of lagged base-year income, plus the deviation between the log of base-year income and the log of lagged base-year income; and controlling for a third-order polynomial in the log of lagged base-year income plus a set of demographic characteristics.

Similar to the findings of Kleven and Schultz (2014), the table shows that what matters is that we control for (lagged) base-year income, not how it is controlled for. Whether we control linearly or more flexibly does not materially change the estimates. Once we include a function of lagged base-year income in equation (3.2), the estimated ETIs are all within the range of 0.11 to 0.16, with our baseline estimate of 0.13 in the middle of this range.

Table 3.2: ETI estimates – baseline specification and alternative controls for differential trends

	Baseline	Alternative controls for differential trends					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ETI	0.1346	-0.0001	0.1103	0.1346	0.1097	0.1607	0.1165
	(0.0090)	(0.0085)	(0.0088)	(0.0090)	(0.0093)	(0.0091)	(0.0089)
Controls:							
No income controls		x					
Log income			x				
Third-order polynomial in log income	x			x			
Splines of log income					x		
Third-order polynomial + deviation						x	
Third-order polynomial + demographic controls							x
First stage coefficient	0.2998	0.3143	0.3064	0.2998	0.2918	0.2986	0.3002
	(0.0007)	(0.0007)	(0.0007)	(0.0007)	(0.0008)	(0.0007)	(0.0007)
First stage F-statistic	165,358	185,331	172,032	165,358	144,681	164,138	166,963
Observations	18,707,745	18,707,745	18,707,745	18,707,745	18,707,745	18,707,745	18,707,745
Individuals	2,600,580	2,600,580	2,600,580	2,600,580	2,600,580	2,600,580	2,600,580

Notes: This table reports results from estimating equation (3.2) using 2SLS, instrumenting the change in the net-of-tax rate with the mechanical change in the net-of-tax rate defined in equation (3.3), with different approaches to accounting for differential income growth. Regressions are weighted by lagged base-year income and estimated on stacked panel datasets including three-year differences covering the period 1999–2018, except the years 2010–2013. "Log income" refers to the log of lagged base-year income, "splines of log income" refers to a ten-piece spline in the log of lagged base-year income, and "deviation" refers to the deviation between the log of base-year income and the log of lagged base-year income. Demographic controls include controls for age (second-order polynomial), sex, marital status, educational attainment, and municipality of residence in a densely populated region. Standard errors are reported in parentheses and clustered by individuals.

3.3.2 Robustness checks

In Table 3.3, we show results from three different robustness checks: estimating equation (3.2) without income weights; excluding the bottom 20% of the income distribution; and defining the mechanical change in the net-of-tax rate (the instrument) the same way as in Gruber and Saez (2002).[10] Omitting the income weights and excluding the bottom 20% of the income distribution results in ETI estimates that are very close to our baseline estimate of 0.13, while using the Gruber and Saez (2002) instrument gives a much larger first-stage coefficient and a somewhat smaller ETI estimate.

Table 3.3: Robustness checks – regression weights, sample restrictions, and instrument definition

	(1)	(2)	(3)	(4)
ETI	0.1346	0.1504	0.1361	0.0868
	(0.0090)	(0.0078)	(0.0113)	(0.0043)
Baseline specification	x			
No income weights		x		
Exclude bottom 20%			x	
Gruber and Saez (2002) instrument				x
First stage coefficient	0.2998	0.2575	0.2974	0.6461
	(0.0007)	(0.0006)	(0.0008)	(0.0007)
First stage F-statistic	165,358	176,300	125,337	945,401
Observations	18,707,745	18,707,745	15,327,610	18,707,745
Individuals	2,600,580	2,600,580	2,192,825	2,600,580

Notes: This table reports results from estimating equation (3.2) using 2SLS, controlling for differential income growth by including a third-order polynomial in lagged base-year income and instrumenting the change in the net-of-tax rate with the mechanical change in the net-of-tax rate defined in equation (3.3). Regressions are weighted by lagged base-year income (except for the specification in column (2)) and estimated on stacked panel datasets including three-year differences covering the period 1999–2018, except the years 2010–2013. The Gruber and Saez (2002) instrument is defined according to equation (3.3), but with base-year income y^*_{it-s} replacing the lag of base-year income y^*_{it-s-1}. Standard errors are reported in parentheses and clustered by individuals.

Table 3.4 shows results from estimating equation (3.2) using 2SLS when we vary the interval length s from $s = 1$ up to $s = 5$. In our baseline specification, we use an interval length of $s = 3$, following common practice in the literature. The table shows that an interval length of $s = 1$ gives a somewhat smaller ETI estimate, while the estimated ETIs for s between 2 and 5 are within the range of 0.10 to 0.15, with our baseline estimate of 0.13 in the middle of this range.

3.3.3 ETI estimates for different reform periods

Table 3.5 presents results from our baseline specification estimated on stacked panel datasets including three-year differences covering four different periods: (i) all relevant three-year differences in the period 1999–2018 ("All years"); (ii) the three-year differences 2000–1997, 2001–1998, and 2002–1999 ("2000 reform");

[10] Gruber and Saez (2002) also define the instrument according to equation (3.3), but with base-year income y^*_{it-s} replacing the lag of base-year income y^*_{it-s-1}.

Table 3.4: Robustness checks – alternative interval lengths

	(1)	(2)	(3)	(4)	(5)
ETI	0.0502	0.1053	0.1346	0.1549	0.1322
	(0.0102)	(0.0085)	(0.0090)	(0.0096)	(0.0099)
Interval length s	1	2	3	4	5
First stage coefficient	0.2253	0.2684	0.2998	0.3308	0.3433
	(0.0008)	(0.0007)	(0.0007)	(0.0008)	(0.0008)
First stage F-statistic	79,785	139,795	165,358	165,828	168,213
Observations	23,498,437	20,876,844	18,707,745	16,788,209	15,061,112
Individuals	2,926,299	2,748,151	2,600,580	2,469,774	2,344,834

Notes: This table reports results from estimating equation (3.2) using 2SLS, controlling for differential income growth by including a third-order polynomial in lagged base-year income and instrumenting the change in the net-of-tax rate with the mechanical change in the net-of-tax rate defined in equation (3.3). Regressions are weighted by lagged base-year income and estimated on stacked panel datasets including differences of length s covering the period 1997–2018, except the years 2010–2013. Each column shows results from a specification with a given interval length s, with s varying from 1 in column (1) to 5 in column (5). Standard errors are reported in parentheses and clustered by individuals.

(iii) the three-year differences 2005–2002, 2006–2003, 2007–2004, and 2008–2005 ("2006 reform"); and (iv) the three-year differences 2016–2013, 2017–2014, and 2018–2015 ("2016 reform").

The table shows that the estimated ETIs and first-stage coefficients are positive and statistically significant for all the time periods considered. The ETI estimates for the 2000 and 2006 reforms, with the largest changes in tax rates for high-income earners and around the middle of the income distribution, respectively, are somewhat smaller than our baseline estimate of 0.13, at 0.09 and 0.11. For the 2016 reform, with the largest changes in the lower parts of the distribution, we obtain an ETI estimate of about 0.23.

Table 3.5: ETI estimates for different reform periods

	(1) All years	(2) 2000 reform	(3) 2006 reform	(4) 2016 reform
ETI	0.1346	0.0902	0.109	0.2311
	(0.0090)	(0.0353)	(0.0110)	(0.0450)
First stage coefficient	0.2998	0.2006	0.3624	0.1319
	(0.0007)	(0.0018)	(0.0011)	(0.0019)
First stage F-statistic	165,358	12,728	99,881	4,633
Observations	18,707,745	3,242,025	4,593,232	3,906,827
Individuals	2,600,580	1,337,734	1,518,656	1,609,796

Notes: This table reports results from estimating equation (3.2) using 2SLS, controlling for differential income growth by including a third-order polynomial in lagged base-year income and instrumenting the change in the net-of-tax rate with the mechanical change in the net-of-tax rate defined in equation (3.3). Regressions are weighted by lagged base-year income and estimated on stacked panel datasets including three-year differences covering four different periods: (i) all three-year differences in the period 1999–2018, except the years 2010–2013 ("All years"); (ii) the three-year differences 2000–1997, 2001–1998, and 2002–1999 ("2000 reform"); (iii) the three-year differences 2005–2002, 2006–2003, 2007–2004, and 2008–2005 ("2006 reform"); and (iv) the three-year differences 2016–2013, 2017–2014, and 2018–2015 ("2016 reform"). Standard errors are reported in parentheses and clustered by individuals.

4 Earnings responses to lotteries

4.1 Data, sample, and representativeness

Our empirical analyses combine several administrative data sources that we link together using unique identifiers for individuals and households. This results in a panel dataset covering the full Norwegian population in the period 1995–2018. In addition to individual characteristics such as age, sex, and marital status, the dataset includes detailed information from tax records on income and wealth, including lottery winnings[11] and earnings, as well as the market values of most assets. We construct measures of each household's net-worth following the methodology of Eika et al. (2020): We supplement the tax records on wealth with measures of market values of real estate, using data on transactions in real estate, information on the characteristics of each property, and detailed housing price indices.

Although Lottery winnings are exempt from income tax, individuals are legally required to report winnings of NOK 100,000 and more to the tax authorities.[12] We measure the size of the win on a per-adult basis by normalizing it by the number of adults in the household (one if single and two if married or registered as cohabitants). This normalization reflects that most lottery winnings accrued during marriage are de facto treated as being owned by spouses equally. Earnings are defined as the the sum of wage earnings and self-

[11]In Norway, two state-owned companies are allowed to offer gambling services: Norsk Tipping (mainly lotteries and betting on sports events) and Norsk Rikstoto (horse racing). With around 2 million players in every year, lottery participation is by far the most popular gambling activity in Norway (see `https://2020.norsk-tipping.no/en/the-gaming-market/`). We therefore refer to winners and winnings from any gambling activity as lottery winners and winners.

[12]More precisely, since 2007 individuals are required to report winnings of NOK 100,000 or more. Before 2007, the reporting threshold was NOK 10,000.

employment income of the winner, and someone is defined to be employed if her earnings exceed one basic amount in a given year.

To construct our main estimation sample, we impose three restrictions. First, we focus on winners who won a prize of at least NOK 100,000 once during our sample period. We choose this amount since any win of this size must be reported to the tax authorities in all years in our sample. To separate gambling wins that resemble income from an entrepreneurial activity (e.g., professional poker) from windfall gains due to lottery winnings, we also exclude repeated winners (with prizes above the reporting threshold) from our sample. Second, we require each individual to be between age 25 and 61 in their win year since our main focus is on labor supply responses of working-age individuals. Lastly, we require each individual to be in the sample for at least two years prior to winning. The resulting sample consists of more than 14,000 unique winners with an average win of around $80,000 across 22 win-year cohorts.

Table 4.1: Summary statistics of individual characteristics and labor market outcomes

	Winners (Age 25 - 61)	Population (Age 25 - 61)
Earnings	54,402	49,129
Employment	0.88	0.83
Net-worth	132,249	125,242
Age	45.46	42.99
Number of children	1.03	1.10
Years of education	12.60	12.76
Male	0.57	0.50
Married	0.55	0.55
Q1 share	0.19	0.25
Q2 share	0.24	0.25
Q3 share	0.27	0.25
Q4 share	0.30	0.25

Notes: This table presents a summary of the descriptive statistics in our sample of working-age winners. All monetary values are adjusted for inflation using the Consumer Price Index and reported in 2018 U.S. dollars using an exchange rate of 8.13 NOK/USD. In the first section of the table, we report mean characteristics. All values for the winners sample are measured in the year prior to the win year and reported as cohort-size-weighted averages. The final column reports the same set of descriptive statistics for the population aged 25-61, taking a cohort-size weighted average across the 1996 to 2017 years. Net-worth is measured at the household level and normalized by the number of adults living in the household. In the second section of the table, we present a comparison of the pre-win distribution of winner earnings to that in the population aged 25-61. For each calendar year, we map each winner to the corresponding quartile in the earnings distribution of the working-age population. We then calculate the share of winners falling into each quartile of this distribution. Finally, we take the mean of the shares across calendar time (for each quartile). For the working-age population, this share is mechanically 0.25.

Column 1 in Table 4.1 reports a set of key summary statistics for our sample. All summary statistics for the winners are measured in their baseline year, i.e., one year prior to their win year. Each statistic is calculated as a weighted average using cohort size as weights. In column 2, we compare winners to the average working-age population. Winners are more likely to be male, and have somewhat higher earnings and employment rates, but also tend to be slightly older compared to the average working-age resident. By and large, however, we find that lottery winners are broadly similar to the working-age population.

The last four rows of Table 4.1 show how the income distribution of winners compares to that of the working-age population. We do so by calculating the share of individuals in our sample that falls into each earnings quartile in the working-age population. Overall we find that winners are well represented in each earnings quartile.

In our investigation of heterogeneous responses to winning the lottery, we group winners according to their average earnings in the two years preceding the lottery win. To better align the winners sample with the analyses of the specific reforms in the top, middle and lower income brackets in Section 3, we drop from our sample individuals in the first quartile of the pre-win earnings distribution. Individuals in the second quartile are individuals with positive earnings who are primarily located in the lowest income tax bracket, and we thus refer to them as low-income earners. Individuals in the third and fourth quartile are individuals with positive earnings who are primarily located in the middle and top income tax bracket, respectively. We therefore refer to these groups as medium and high-income earners.

4.2 Event-study design and average effects of winning

We now present and apply the event-study design that we use to draw causal inferences about how individuals adjust their labor supply in response to winning the lottery. This design exploits variation in the timing of lottery wins, and, thus, can be expressed as a difference-in-differences (DiD) estimator that compares the early winners to those who win later.

To fix ideas, consider a specific cohort of winners who win in year w. Our parameter of interest is the cohort-specific average effect of the lottery win on some outcome y as measured in post-win year $w + \ell$. The key identification challenge is that we do not observe the counterfactual outcome of winners had they not won. A natural approach to recover our parameter of interest is to construct a DiD estimator using a control group that would arguably have experienced the same change in outcomes from before to after the lottery win in the absence of winning. One natural candidate for such a control group are individuals who win in later years. The DiD estimator for cohort w then is

$$\underbrace{\mathbb{E}\left[y_{i,w+\ell} - y_{i,w-1} \mid i \text{ won in } w\right]}_{\text{difference over time for treatment group}} - \underbrace{\mathbb{E}\left[y_{i,w+\ell} - y_{i,w-1} \mid i \text{ has not won by } w+\ell\right]}_{\text{difference over time for control group}} . \tag{4.1}$$

The DiD estimator eliminates unobserved time-invariant individual heterogeneity by comparing winners before and after they win, while accounting for year and event-time effects by using the later winners as a control group before they win. As long as the individuals in the treatment and control groups would have had a common trend between years $w - 1$ and $w + \ell$ in the expected outcome in the absence of winning the lottery, the DiD estimator defined in expression (4.1) recovers the average impact of lottery winning for cohort w in year $w + \ell$ for $\ell \geq 0$.

To implement the DiD estimator defined in expression (4.1), we use a regression to make it easier to include additional covariates and calculate standard errors. For each cohort w and each event time ℓ we create a subsample of the treated individuals who won in period w and a control group of individuals who have not won by period w or $w + \ell$, whichever is greater. Using this subsample, we run the regression

$$y_{i,t} = \alpha_1^{w,\ell} + \alpha_2^{w,\ell} \mathbb{1}\{i \text{ won in } w\} + \alpha_3^{w,\ell} \mathbb{1}\{t = w + \ell\} + \rho^{w,\ell} z_{i,t} + u_{i,t}^{w,\ell}, \tag{4.2}$$

where $z_{i,t}$ represents the interaction term $\mathbb{1}\{i \text{ won in } w\} \times \mathbb{1}\{t = w + \ell\}$. Here, $\alpha_1^{w,\ell}$ is the control group mean in the year before the lottery win, $\alpha_2^{w,\ell}$ is a fixed effect for the treated individuals (cohort w), $\alpha_3^{w,\ell}$ is a time effect for event time ℓ, and $\rho^{w,\ell}$ is an interaction effect and our parameter of interest. We estimate the model separately for each cohort w and then take a weighted average of the estimates for each event time ℓ, where the weights are determined by the cohort-size.[13]

In Figure 4.1 we plot the estimated coefficients from this regression for the full sample. For each outcome and each event time ℓ, we report a cohort-weighted average of regression coefficient $\rho^{w,\ell}$, with the baseline event time $w - 1$ normalized to zero. There is no evidence of systematically different time trends between current and later winners in pre-win event times -5 to -1 for any of the outcomes. This is consistent with the outcomes of current winners and the later winners evolving in the same way across years in the pre-win period, providing support for the common trends assumption.

Figure 4.1a shows that winning the lottery leads to a significant and immediate increase in net-worth in the win-year. This increase in net-worth relative to later winners then declines in the years thereafter. This pattern is consistent with lottery winners first expanding their savings and then gradually consuming out of their new wealth. If leisure is a normal good, then economic theory predicts that an unanticipated increase in wealth leads to a reduction in lifetime labor supply. Figure 4.1b shows that winning the lottery indeed leads to a sizable, swift and persistent reduction in earnings of the winner relative to later winners. On average, earnings of the winner fall by around $2,800 (approximately 4%) in each period following the lottery win. It is also evident from Figure 4.1c that winners are more likely to stop working, which accounts for a non-negligible share of the reduction in earnings that we observe.

These average effects of winning on earnings and employment mask considerable heterogeneity across the earnings distribution. In Figure 4.2 we explore this heterogeneity in responses by splitting the sample according to the average earnings in the two years preceding the lottery win. As one can see from these figures, high-income earners reduce their earnings compared to low-income earners by more than twice as much. In contrast, low- and medium-income households, however, are considerably more likely to stop working. Their employment rates have decreased on average by around 2 percentage points within a few years after winning the lottery. Given the systematic heterogeneity across the earnings distribution, we will hereafter look separately at the impacts for low-, medium- and high-income earners.

[13] To control for age composition, we include a full set of dummies for each age in regression model (4.2). These controls adjust for the fact that current winners are slightly older than later winners in year w, which is to be expected if the timing of win is as good as random. For other observables, current and later winners have very similar pre-win characteristics.

Figure 4.1: Effects of winning on wealth, earnings and employment

(a) Net-worth

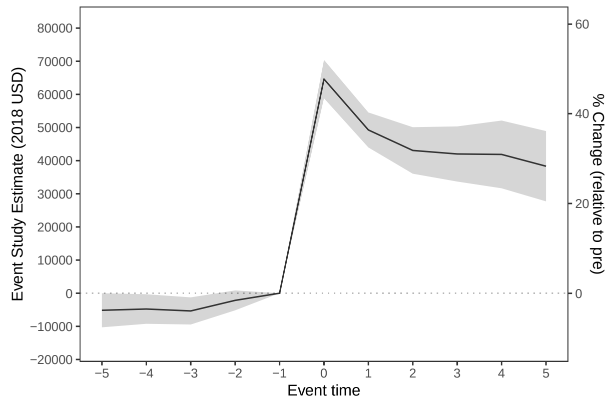

(b) Earnings

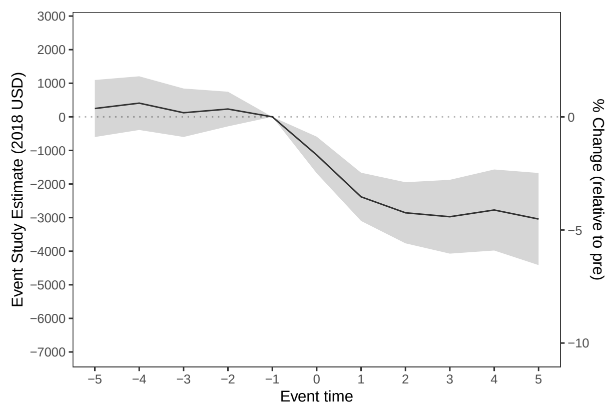

(c) Employment

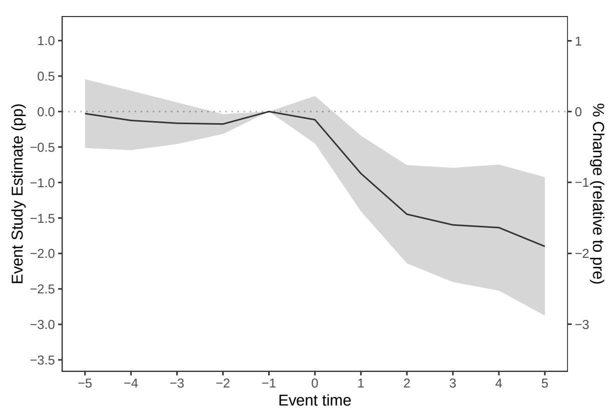

Notes: This figure presents estimates of the impact of winning on net-worth, earnings and employment based on estimating a version of equation (4.2) for each outcome, and then taking cohort-size-weighted averages of for each event time. 90 percent confidence intervals are displayed, clustering on winner. Throughout, we use $w-1$ as the omitted event time. In addition to the cohort-size-weighted average effect in levels (left-hand axis), each subfigure also reports this average effect scaled by the mean of the outcome in omitted event time (right-hand axis) which can be interpreted as an average percentage change (relative to the pre-win period) in the outcome.

Figure 4.2: Effects of winning on wealth, earnings and employment

(a) Low-income earner
Earnings

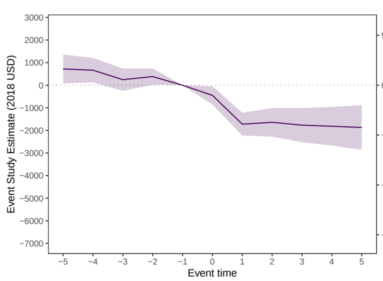

(b) Medium-income earner
Earnings

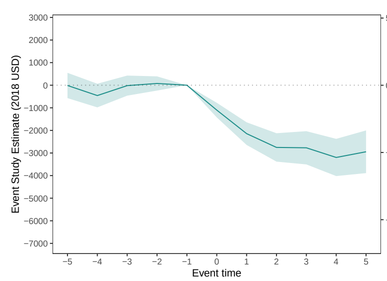

(c) High-income earner
Earnings

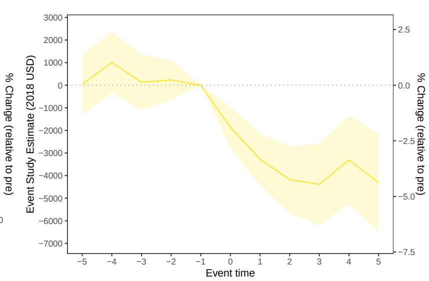

(d) Low-income earner
Employment

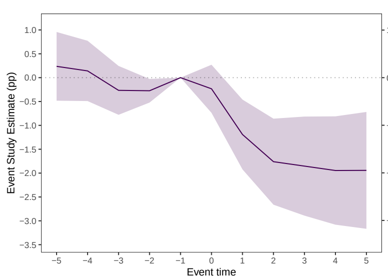

(e) Medium-income earner
Employment

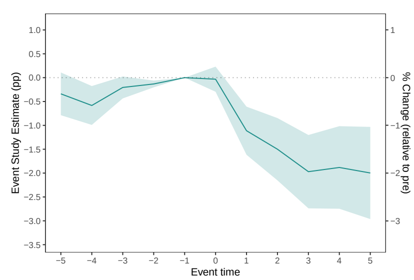

(f) High-income earner
Employment

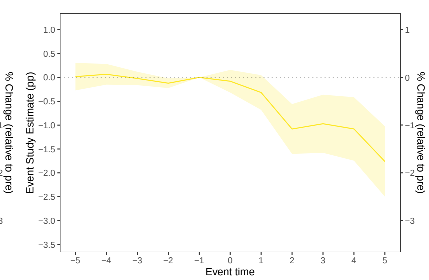

Notes: This figure presents estimates of the impact of winning on earnings and employment based on estimating a version of equation (4.2) separately by income groups for each outcome, and then taking cohort-size-weighted averages of for each event time. 90 percent confidence intervals are displayed, clustering on winner. Throughout, we use $w - 1$ as the omitted event time. In addition to the cohort-size-weighted average effect in levels (left-hand axis), each subfigure also reports this average effect scaled by the mean of the outcome in omitted event time (right-hand axis) which can be interpreted as an average percentage change (relative to the pre-win period) in the outcome.

4.3 Earnings responses to increases in wealth and unearned income

Estimates of wealth effects

It is difficult to interpret the size of the effects reported in the previous subsection because the treatment variable captures whether a person wins but not the size of the win. To get economically-interpretable estimates, we now shift to an IV model that uses variation in the timing of the lottery wins as an instrument for lottery winnings. The resulting IV estimates tell us individual responses per dollar of winnings, which we will refer to as wealth effects.

We maintain the conventions and notation from our event-study regression model (4.2) in the prior subsection. For each cohort w and each event time $\ell \geq 0$, we estimate the following IV model:

$$x_{i,t} = \mu_1^{w,\ell} + \mu_2^{w,\ell}\mathbb{1}\{i \text{ won in } w\} + \mu_3^{w,\ell}\mathbb{1}\{t = w + \ell\} + \phi^{w,\ell} z_{i,t} + \varepsilon_{i,t}^{w,\ell} \tag{4.3}$$

$$y_{i,t} = \theta_1^{w,\ell} + \theta_2^{w,\ell}\mathbb{1}\{i \text{ won in } w\} + \theta_3^{w,\ell}\mathbb{1}\{t = w + \ell\} + \beta^{w,\ell} x_{i,t} + \nu_{i,t}^{w,\ell} \tag{4.4}$$

Starting with the first-stage equation (4.3), the parameters $\{\mu_1^{w,\ell}, \mu_2^{w,\ell}, \mu_3^{w,\ell}\}$ capture cohort and time effects. The endogenous variable x is the lottery win in the win year. Thus, the first-stage coefficient $\phi^{w,\ell}$, which captures the impact of a lottery win on lottery winnings, does not change over time. The second-stage equation (4.4) relates our outcome of interest y to changes in x. The reduced form of the IV model is given by the event-study regression (4.2).

Table 4.2 reports estimates of the average annual earnings and employment response to an additional dollar of wealth separately for low-, medium- and high-income earners. We estimate that for an extra 100 dollars in wealth, winners reduce their earnings on average by around 4 dollars both in the short- and long-run. The similarity in the estimates suggests that a windfall gain has a persistent impact on labor supply. Across the income distribution, we observe that the average earnings reduction per dollar of additional wealth is increasing in pre-win income. For example, whereas low-income individuals reduce their own annual earnings by 2.5 dollars per 100 dollars of additional wealth, higher-income individuals decrease their annual wage earnings by over twice as much. The lower- and medium-income individuals, however, are more likely to stop working – the reduction in the probability of employment for them is almost 1 percentage point larger than for winners with higher income.

These sizable and persistent earnings and employment responses raise a natural question: How much of the overall earnings response is attributable to the extensive margin? To address this question, we decompose the earnings response into extensive- and intensive-margin contributions. Concretely, we take a standard statistical intensive-extensive decomposition of cross-sectional earnings effects and adapt it to our DiD estimator (4.1). Table 4.2 shows the share of the observed earnings response that is attributable to the extensive-margin response. On average, we find that the extensive margin explains roughly half of the earnings response in the short-run, and close to 60% in the longer-run. The importance of the extensive margin, however, decreases with pre-win income. For example, whereas the extensive margin explains 56% of the observed earnings response for low-income individuals in the short-run, 41% of the response is explained by employment responses for higher-income individuals.

Table 4.2: Wealth effects on earnings and employment

	Low-income earners	Medium-income earners	High-income earners	Average
		Short-run		
Earnings	-2.5596	-3.7677	-5.8417	-4.0563
(per $100)	(0.6211)	(0.5888)	(1.3986)	
Employment	-0.0261	-0.0282	-0.0156	-0.0233
(per $100,000)	(0.0091)	(0.0067)	(0.0055)	
Extensive-margin share	56%	49%	41%	49%
		Longer-run		
Earnings	-2.4837	-4.1765	-5.9695	-4.2099
(per $100)	(0.7458)	(0.7395)	(1.4281)	
Employment	-0.0284	-0.0317	-0.0253	-0.0285
(per $100,000)	(0.0103)	(0.0086)	(0.007)	
Extensive-margin share	65%	57%	56%	59%

Notes: This table presents estimates of the mean effect of an extra dollar of wealth on earnings and employment of the winner. These estimates are calculated by first estimating the 2SLS regression, defined by equations (4.3) and (4.4), for each outcome, and then taking cohort-size weighted averages of $\beta^{w,\ell}$ for each event time ℓ. The short-run estimates refer to event time $\ell = 3$, while the longer-run estimates refer to event time $\ell = 5$. We report wealth effects for each subsample of winners with low, medium and high average incomes in the two years prior to winning. In addition, we also report the share of the observed earnings responses that is attributable to the extensive margin.

Estimates of marginal propensities to earn

The size of the wealth effects reported in Table 4.2 can be hard to gauge as the observed responses to windfall gains should vary across individuals depending on a number of factors, such as the age at which the individual wins, her savings behavior, and the tax rates she faces. Guided by the theory outlined in Section 2, we are ultimately interested in the share of per-period unearned income that is spent on reducing labor supply. In other words, we are interested in the marginal propensity to earn (MPE) governing labor supply behavior.

As a preliminary step, we therefore need to infer the effect of lottery winnings on unearned income. One popular and simple approach to determine how windfall gains are allocated across time is the annuitization method. If we assume that winnings are smoothed perfectly over the remaining lifetime, then it is easy to show that the unearned income of a k-year old winner must increase by

$$\frac{r}{1+r}\left(1-\left(\frac{1}{1+r}\right)^{T-k+1}\right)^{-1}L, \tag{4.5}$$

where r is the post-tax interest rate, L is the amount won, and $T-k$ is the remaining lifetime of the winner. In our empirical analysis, we assume that all individuals live for $T=80$ years, and set $r=3\%$, which is close to the average yearly interest rate on deposits in Norway for our period of observation.[14] An alternative approach where one does not make any assumptions a priori about household behavior relies on rich administrative records on income and wealth. In this case one can use this information to directly approximate unearned income at the individual level.

Using the IV model defined by equations (4.3) and (4.4), we can estimate how an extra dollar in unearned income translates into a decrease in earnings. This is done by 2SLS estimation of the two-equation system with the endogenous variable being the unearned income in a given period, and the outcome variable being winner earnings. The resulting estimates then capture both intensive- and extensive-margin responses. For the analyses in Section 5, however, we need to isolate the MPEs along the intensive-margin, which we can directly infer from these estimates based on the extensive-margin share reported in Table 4.2.

Figure 4.3 presents the intensive-margin MPEs for the full sample and separately by income groups. The estimates here are reported using the annuitization method, but on average we find very similar, albeit less precise, estimates when we use unearned income as observed in the data (see Appendix Figure A.1). Figure 4.3 shows that earnings responses to a change in unearned income along the intensive margin are quite large and fairly stable over time. An extra dollar in unearned income leads approximately to a 47 cent reduction in earnings on average. Furthermore, there is substantial heterogeneity in intensive-margin MPEs across the income distribution. For lower-income individuals, the intensive-margin earnings response is around 24 cents per dollar. For those with higher incomes, the response is around 74 cents per dollar and hence more than three times as large.

[14] See https://www.ssb.no/en/statbank/table/08175/tableViewLayout1/

Figure 4.3: Intensive-margin MPEs

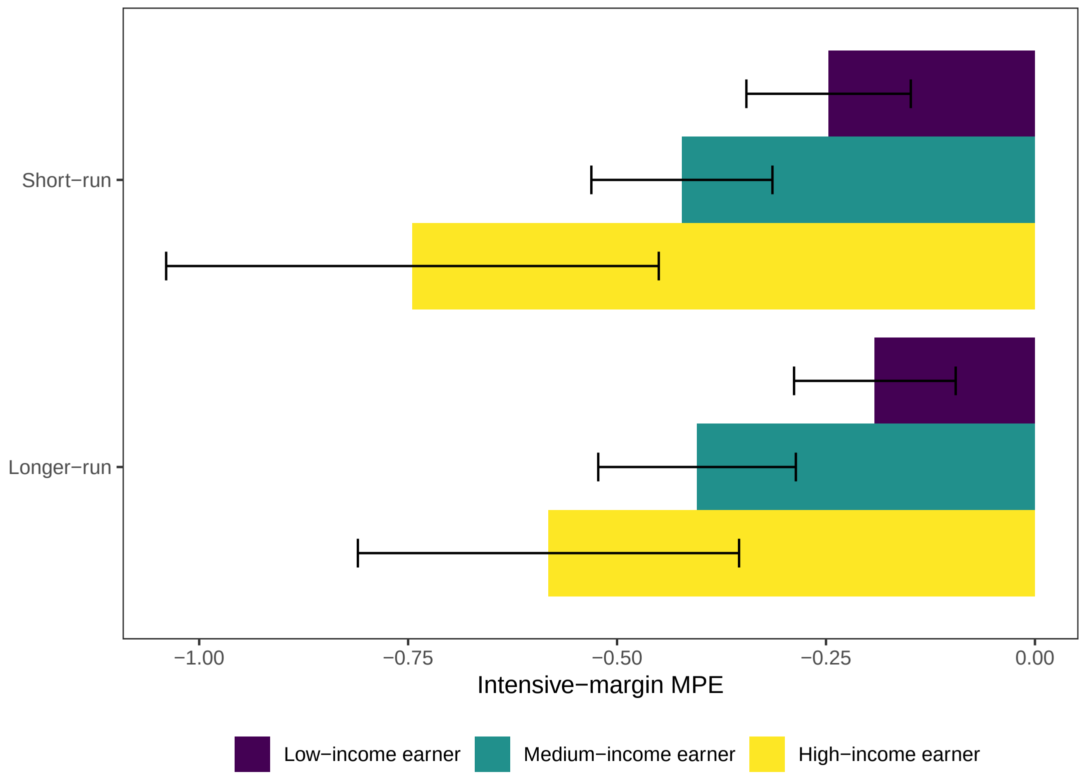

Notes: This figure presents estimates of the intensive-margin MPEs. These estimates are calculated by first estimating the 2SLS regression, defined by equations (4.3) and (4.4), using unearned income as the endogeneous variable. We then take cohort-size weighted averages of $\beta^{w,\ell}$ for each event time ℓ. The short-run estimates refer to event time $\ell = 3$, while the longer-run estimates refer to event time $\ell = 5$. To calculate the intensive-margin MPE, we subtract the component that is due to responses along the extensive margin, using the extensive-margin shares reported in Table 4.2. We report intensive-margin MPEs for each subsample of winners with low, medium, and high average incomes in the two years prior to winning. 90 percent confidence intervals are displayed, clustering on winner.

5 Recovering elasticities and policy analysis

In this section, we combine the model developed in Section 2 with the empirical estimates of earnings responses to lottery winnings and tax reforms to point identify or bound the Slutsky parameters. We next use these estimates to assess the deadweight loss of taxation and the optimal marginal tax rates.

5.1 Recovering model elasticities

Motivated and guided by the empirical results discussed above, we separately consider model elasticities for low, middle and high income earners. These groups are denoted L, M and H. We restrict the intensive margin elasticities to be homogeneous within groups but potentially heterogeneous across groups – a weaker assumption than the usual full homogeneity assumption in the existing empirical literature.

How to recover elasticities in the top income tax bracket

Consider first the simplest case of the top income tax bracket. $H-$households have income $y_i \geq \underline{y}_H$ and face a marginal tax rate of τ_H. Using equation (2.12), the ETI in the $H-$bracket can be expressed as:

$$\mathbb{E}_{\geq \underline{y}_H} \text{ETI}_i = \text{ETI}_H = \left[\zeta_H^c + \mathbb{E}_{\geq \underline{y}_H} \left(\frac{y_i - \underline{y}_H}{y_i} \right) (1 - \tau_H) \eta_H, \right] \tag{5.1}$$

where ζ_H^c and η_H are the compensated elasticity and income effect for $H-$households. Using equations (2.8) and (2.10), we can express the average earnings response to an increase in unearned income in the top bracket, as

$$\mathbb{E}_{\geq \underline{y}_H} \text{MPE}_i = \text{MPE}_H = \left[\frac{\eta_H}{1 + \theta \zeta_H^c} \right] - \int_{\underline{y}_H}^{\infty} y_i \lambda_L(y_i) dH(y), \tag{5.2}$$

where the last term, the integral, captures the extensive margin response. From the data, we can estimate $\mathbb{E}_{\geq \underline{y}_H} \left(\frac{y_i - \underline{y}_H}{y_i} \right)$ and $\int_{\underline{y}_H}^{\infty} y_i \lambda_L(y_i) dH(y)$. Applying these numbers and the estimated earnings responses from the tax reforms and the lotteries for $H-$households, we obtain the following two equations with two unknowns, which can be solved to point identify the parameters of interest:

$$0.09 = \zeta_H^c + 0.22(1 - 0.472)\eta_H \tag{5.3}$$

$$-0.74 = \left[\frac{\eta_H}{1 + 0.15 \zeta_H^c} \right]. \tag{5.4}$$

It is useful to observe that to get point identification of the income effect and the compensated elasticity (and by implication, the uncompensated elasticity), it is essential to know the earnings responses both to the lottery winnings and to the tax reforms. With information about only one of the data moments, say we only knew $\text{ETI}_H = 0.09$, we could only - at best - bound the Slutsky parameters. If leisure is a normal good, we know that $\eta_H \leq 0$ and with standard convex preferences we know that the compensated elasticity is weakly positive. Thus, $\zeta_H^u \leq 0.09, \zeta_H^c \geq 0.09$ and $\eta_H \leq 0$. Likewise, provided only with results from lottery responses we could use this information to bound the Slutsky parameters. The intensive margin earnings response to a one dollar increase in unearned income is estimated to be -0.74. With a non-linear tax system, this response constitutes an upper bound of the income effect (see equation 5.4), hence $-0.74 \geq \eta_H$. Convex preferences give $\zeta_H^c \geq 0$. Note, however, that the bounds on the income parameter and on the compensated elasticity put no constraints on the magnitude of the uncompensated elasticity.

How to recover elasticities in brackets other than the top

In a similar fashion, we can use the model and data to either bound or point identify the income effects and compensated elasticity for the L and M households. However, one complication arises in the mapping between the estimates of ETI for L and M and the model parameters. The issue is that an increase (decrease) in the marginal tax rate in a lower bracket increases (reduces) the average tax rate in higher brackets. To illustrate, consider a reform that increases the net of tax rate in bracket M by $\triangle$. Although this reform has no "price" effect in bracket H, it reduces the average tax rate for $H-$households and, thus, increases their income. Their response to this income effect is to reduce earnings by $\triangle \left[\underline{y}_H - \underline{y}_M\right] \eta_H$. Since this is a response to the reform that lowers post-reform earnings in the control group in the ETI estimation, it creates an upward bias in the causal effect of an increase in the net of tax rate for $M-$earners. Once we have made this adjustment of the ETI estimates, we can follow the same identification procedure as we did for the H-households in equations 5.3 and 5.4.

Empirical results

Table 5.1 presents bounds and point estimates for all three income brackets. For L and M the numbers are adjusted for the potential bias discussed above.[15] The point estimates show that income effects are considerable, especially for high income earners who reduce their earnings by 76 cents per dollar increase in unearned income. The compensated elasticities range from 0.28 for the low-income earners to 0.15 for medium-income earners. The uncompensated elasticities never exceed 0.23 in absolute value.

5.2 Implications for top tax rates

There is a long-standing debate about how hard high income earners should be taxed. While redistribution through taxes and transfers can increase social welfare by making the distribution of income more equal, it comes at the cost of distorting work incentives and, consequently, decreased economic efficiency. This is the classic trade-off between efficiency and equity which is at the core of the optimal income taxation problem. In a large class of models, this trade-off largely depends on the earnings elasticities, the shape of the income distribution, and the redistributive goals of the government. In this section, we explore what our empirical findings imply for efficiency losses and optimality of top tax rates.

In principle, it is also possible to analyze the efficiency losses and optimality of tax rates in brackets other than the top. However, this would require either strong assumptions about or additional empirical analysis of extensive margin responses to changes in the relevant marginal tax rates.

[15] For the the M group, this adjustment turns out to have a negligible impact on the ETI estimates, mainly because $H-$earners constitute less than 4% of the control group. The adjustment matters slightly more for the estimation of the $L-$reform, where the bias correction reduces the ETI from 0.23 to 0.22.

Table 5.1: Bounds and point estimates of the Slutsky parameters

		ETI	Lottery	ETI & Lottery
High-income earner	ζ_H^u	≤ 0.09		$= -0.23$
	ζ_H^c	≥ 0.09	≥ 0	$= 0.17$
	η_H	≤ 0	≤ -0.74	$= -0.76$
Medium-income earner	ζ_M^u	≤ 0.11		$= -0.09$
	ζ_M^c	≥ 0.11	≥ 0	$= 0.15$
	η_M	≤ 0	≤ -0.42	$= -0.43$
Low-income earner	ζ_L^u	≤ 0.22		$= 0.11$
	ζ_L^c	≥ 0.22	≥ 0	$= 0.28$
	η_L	≤ 0	≤ -0.24	$= -0.25$

Notes: This table presents the bounds and point estimates of the Slutsky parameters.

Efficiency loss

As described in Section 2, there are different ways to measure efficiency losses from income taxation. We focus on the normalized deadweight loss from a slightly higher top marginal tax rate. Assuming homogeneous income elasticities among high income earners, and using equation (2.23), we can express the average marginal deadweight loss per unit of additional tax revenues raised as

$$\frac{\mathbb{E}_{\geq \underline{y}_H} \text{MDWL}_i}{\mathbb{E}_{\geq \underline{y}_H} \text{MTR}_i} = \frac{\text{MDWL}_H}{\text{MTR}_H} = \frac{\frac{\tau_H}{1-\tau_H} \zeta_H^c a}{a - 1 + \tau_H \left[(1-a)\, \eta_H - a \frac{\zeta_H^c}{1-\tau_H} \right]}. \tag{5.5}$$

Equation (5.5) expresses the efficiency costs of taxation in terms of three parameters: the compensated elasticity, the income effect, and the parameter $a \equiv \mathbb{E}_{\geq \underline{y}_H} \frac{y_i}{\underline{y}_H}$ which depends on the shape of the income distribution.

In 2015, the top marginal tax rate was 47.2% (excluding employers' national insurance contributions) and we calculate a to be 1.43 using tax returns data from that same year. Together with the estimates from Column 3 of Table 5.1, we can then evaluate expression (5.5) and obtain a normalized efficiency loss of approximately 0.59. This means that for an extra dollar of taxes raised, the government imposes an additional cost of 59 cents on taxpayers in the top bracket (on top of the dollar paid in taxes).

The parameter a is a measure of the thickness of the tail of the income distribution and it is useful to examine it more closely to understand the sizable estimate of deadweight loss. In Figure 5.1, we plot the value of a as a function of income using 2015 tax return data. At the cut-off for the top income tax bracket (vertical dashed line), a takes a value of around 1.43. We also see that the ratio is fairly stable over the right tail of the income distribution which suggests that the top tail can be well approximated by a Pareto distribution. It further implies that a shift in the top tax bracket to the left or right would yield a similar efficiency loss all else equal.

A natural question is how the thickness of the tail of the income distribution affects the deadweight loss

Figure 5.1: The tail of the income distribution

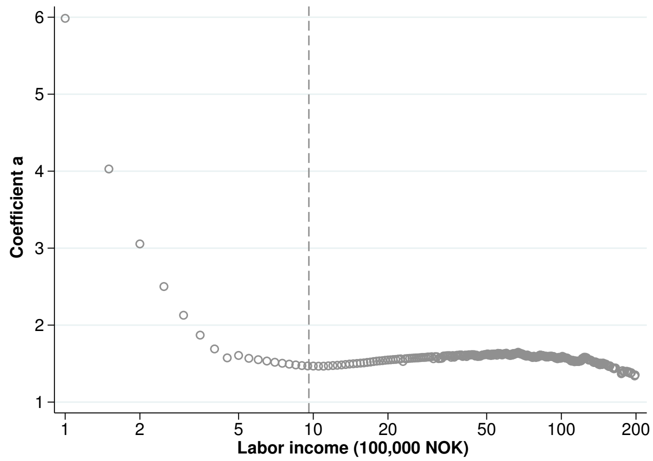

Notes: This figure shows the ratio $a(y) \equiv \mathbb{E}_{\geq y} \frac{y_i}{y}$ using 2015 tax return data. Labor income is displayed on a logarithmic scale and measured in 2018 NOK. The vertical dashed line represents the bracket cut-off for the top tax bracket.

of taxation. A thicker tail leads to higher tax revenues from a marginal change in top tax rates. At the same time it also increases the marginal deadweight loss. It is straightforward to show that the former effect dominates the latter. Thus a thicker tail coincides with smaller efficiency costs of taxation. If we used a coefficient of $a = 3$ in the calculations above, as for example observed in U.S. data[16], we would obtain an efficiency loss of only 20 cents for each additional dollar of tax revenue.

The efficiency costs of taxation (5.5) increase in the compensated elasticity and decrease in the absolute value of the income effect. A careful assessment of the efficiency effects of an income tax reform therefore requires to separate between the two behavioral parameters. In much of the empirical public finance literature the difficulty of separating between the two is often circumvented simply by assuming away income effects. Without income effects the compensated and uncompensated earnings elasticity are the same. Expression (5.1) then implies that these elasticities must be equal to our ETI estimate of 0.09 from Table 5.1. Evaluating expression (5.5) under this assumption leads to a considerably smaller efficiency cost of taxation of approximately 0.37. In other words, assuming no income effects leads one to seriously understate the efficiency loss of taxation.

Lastly, we examine whether we can learn anything about efficiency losses from using responses to lottery winnings or tax reforms alone. As shown in Column 1 of Table 5.1, we can establish a lower bound of 0.09 on the compensated elasticity by using only responses to tax reforms. This implies a lower bound on the marginal deadweight loss which is the numerator in expression (5.5). The denominator in expression (5.5) captures the changes in tax revenues. They can become arbitrarily large as long as income effects are not bounded from below. At the same time, changes in tax revenues can become arbitrarily close to zero if we are to the left of the peak of the Laffer curve. Taken together this implies that using responses to tax reforms alone does not allow us to bound the efficiency losses without further assumptions. Similarly, Column 2 of

[16]See, e.g., Saez et al. (2012) and Golosov et al. (2021).

Table 5.1 shows that we can establish an upper bound of -0.74 on the income effect using only responses to lottery winnings. However, the marginal deadweight loss depends on the compensated elasticity, and without further assumptions it can be arbitrarily large, or small but positive. Everything considered, this means that without further assumptions we cannot bound the efficiency costs from income taxation if we had to rely only on either the observed responses to lottery winnings or tax reforms.

Optimal top tax rates

Optimal income taxation must resolve a trade-off between redistributing income from high income earners to low income earners while minimising the efficiency losses. As discussed in Section 2, the optimal top tax rate τ_H satisfies

$$\frac{\tau_H}{1-\tau_H} = \frac{(1-g_H)(a-1)}{a\zeta_H^c + \tilde{\eta}_H (a-1)}. \tag{5.6}$$

Equation (5.6) expresses the optimal top tax rate in terms of four parameters: the compensated elasticity and the net income effect $\tilde{\eta} \equiv (1-\tau_H)\eta$ governing the behavioral responses to taxes, the parameter a which depends on the shape of the income distribution, and the social marginal welfare weight g_H which reflects the redistributive goals of the government.

The social welfare weight g_H is set by the government. It captures how the government values an additional dollar of consumption for a top bracket earner relative to the tax revenue it can raise. Put differently, g_H is defined such that the government is indifferent between g_H more dollars of tax revenue and one more dollar consumed by a taxpayer in the top income bracket. In Figure 5.2 we plot the optimal top income tax rate (5.6) for different values of g_H. Not surprisingly, less weight on the marginal consumption of high income earners implies higher optimal tax rates. At $g_H = 0$ the government does not put any value on the marginal consumption of high income earners and sets the top rate so as to maximize tax revenue.

Using $a = 1.43$ and the elasticity estimates reported in Column 3 of Table 5.1, we obtain a revenue-maximizing top tax rate of approximately 86%. Conversely, a top income tax rate of 47.2% as observed in 2015 is in line with a social marginal welfare weight of around 0.86. In other words, to justify this level of the top income tax, the decision maker needs to be indifferent between one more dollar of public funds and around 1.16 dollars of additional disposable income to high income earners. If the marginal value of public funds is higher (lower), then the top income tax rate should be increased (decreased).

To better understand the results in Figure 5.2, it is useful to explore how the thickness of the tail in the income distribution affects the optimal top tax rate. It is easy to show that expression (5.6) increases in the thickness of the tail. Intuitively, and as discussed above, a thicker tail coincides with smaller efficiency costs of taxation. This represents a shift in the equity-efficiency trade-off and hence calls for a higher top tax rate to maximize welfare.

The optimal top income tax rate (5.6) is decreasing in the compensated elasticity, and increasing in the absolute value of the income effect. To understand why this is the case, recall that a larger substitution effect implies larger efficiency costs from taxation. Therefore, the tax rate must be lower in optimum. In contrast, a larger absolute value of the income effect leads to a higher optimal tax rate since raising the tax

Figure 5.2: Optimal top tax rate and the government's taste for redistribution

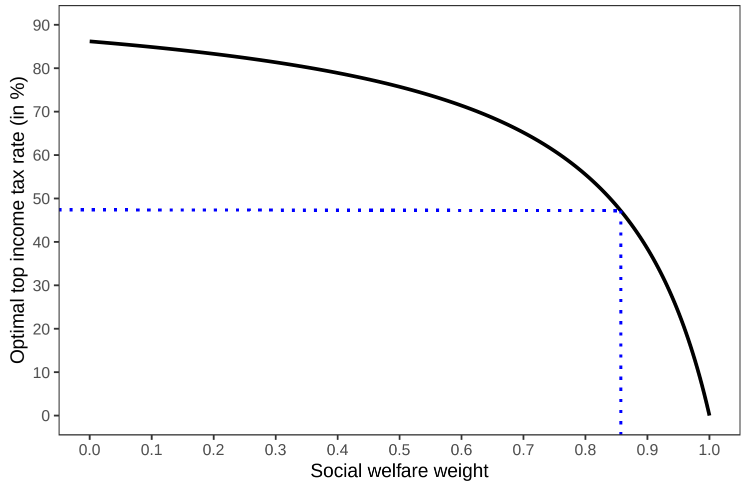

Notes: This figure shows how the optimal top tax rate varies with the social marginal welfare weight of high income earners. Expression (5.6) is evaluated at the value of the compensated and uncompensated elasticity reported in Table 5.1 Column 3. The parameter a is calculated on tax return data from 2015 and set to 1.43. The horizontal dashed line indicates the top tax rate of 47.2% as observed in 2015.

rate reduces disposable income and hence increases labor supply and tax revenue.

We can also ask how our conclusions about optimal top tax rates would change if we assumed that income effects are zero as is typically done in much of the public finance literature. In this case, the compensated and uncompensated earnings elasticity are, by assumption, the same. In this case, the optimal tax formula (5.6) simplifies to $\tau_H = (1-g)(a-1)/\left[\zeta_H a + (1-g)(a-1)\right]$, and expression (5.1) implies that ζ_H must be equal to our ETI estimate of 0.09. This means that both the compensated elasticity and the absolute value of the income effect are lower compared to our baseline parametrization in Column 3 of Table 5.1. A lower compensated elasticity implies lower efficiency costs from taxation which calls for a higher optimal tax rate. Conversely, a zero income effect leads to smaller tax revenues on the margin which in turn implies a lower optimal tax rate. Thus, the overall effect on the optimal tax rate is, a priori, ambiguous. If we keep $g_H = 0.86$ and $a = 1.43$ fixed, and assume away income effects, we obtain a lower optimal top tax rate of around 32% compared to 47% in the baseline parametrization.

References

Auerbach, A. J. and Hines Jr, J. R. (2002). Taxation and economic efficiency. In *Handbook of public economics*, volume 3, pages 1347–1421. Elsevier.

Blomquist, N. S. (1985). Labour supply in a two-period model: The effect of a nonlinear progressive income tax. *The Review of Economic Studies*, 52(3):515.

Blundell, R. and Walker, I. (1986). A Life-Cycle Consistent Empirical Model of Family Labour Supply Using Cross-Section Data. *The Review of Economic Studies*, 53(4):539.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, 1(1):451–488.

Eika, L., Mogstad, M., and Vestad, O. L. (2020). What can we learn about household consumption expenditure from data on income and assets? *Journal of Public Economics*, 189.

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(3):551–572.

Golosov, M., Graber, M., Mogstad, M., and Novgorodsky, D. (2021). How americans respond to idiosyncratic and exogenous changes in household wealth and unearned income. Technical report, National Bureau of Economic Research.

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

Heathcote, J., Storesletten, K., and Violante, G. L. (2017). Optimal Tax Progressivity: An Analytical Framework. *The Quarterly Journal of Economics*, 132(4):1693–1754.

Jacquet, L., Lehmann, E., and Van der Linden, B. (2013). Optimal redistributive taxation with both extensive and intensive responses. *Journal of Economic Theory*, 148(5):1770–1805.

Kleven, H. J. and Kreiner, C. T. (2006). The marginal cost of public funds: Hours of work versus labor force participation. *Journal of Public Economics*, 90(10-11):1955–1973.

Kleven, H. J. and Schultz, E. A. (2014). Estimating Taxable Income Responses Using Danish Tax Reforms. *American Economic Journal: Economic Policy*, 6(4):271–301.

MaCurdy, T. E. (1983). A simple scheme for estimating an intertemporal model of labor supply and consumption in the presence of taxes and uncertainty. *International Economic Review*, 24(2):265–289.

Neisser, C. (2021). The Elasticity of Taxable Income: A Meta-Regression Analysis. *The Economic Journal*, 131(640):3365–3391.

Saez, E. (2001). Using Elasticities to Derive Optimal Income Tax Rates. *The Review of Economic Studies*, 68(1):205–229.

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

Thoresen, T. O. and Vattø, T. E. (2015). Validation of the discrete choice labor supply model by methods of the new tax responsiveness literature. *Labour Economics*, 37:38–53.

Vattø, T. E. (2020). Estimating the elasticity of taxable income when earnings responses are sluggish. *FinanzArchiv*, 76(4):329–369.

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

A Additional Figures

Figure A.1: Comparison between annuitization method and observed intertemporal behavior

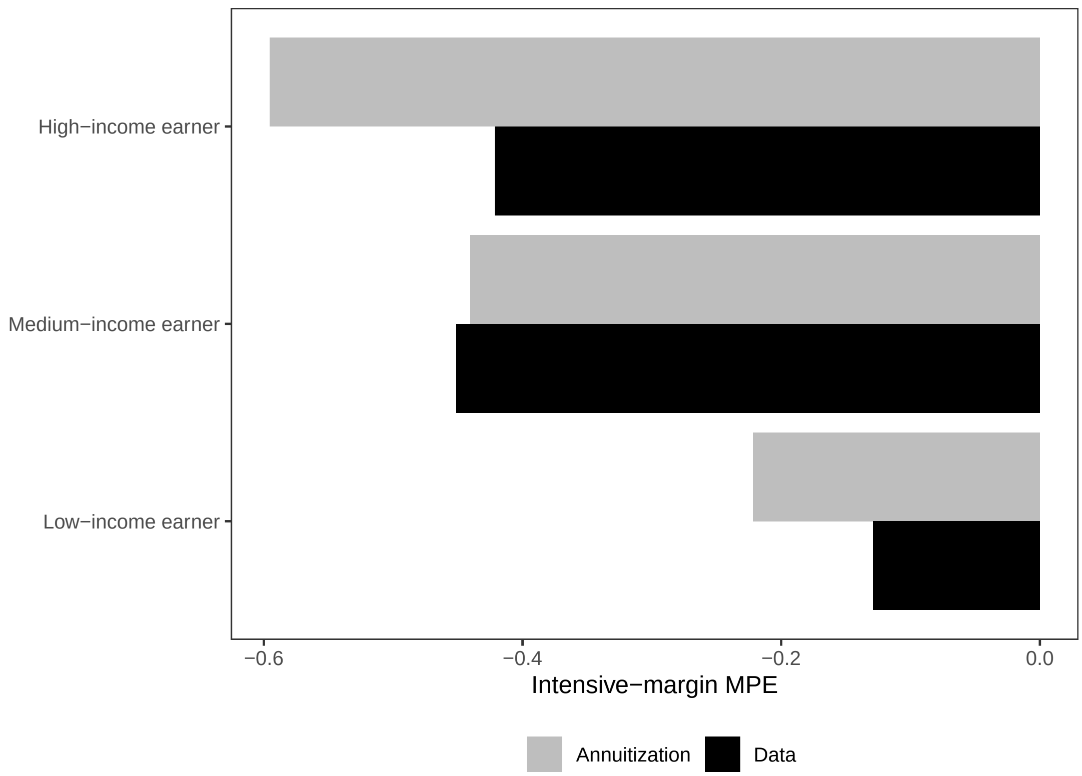

Notes: This figure compares estimates of the intensive-margin MPEs under the annuitization method with the intensive-margin MPEs obtained from actual intertemporal behavior as observed in the data. These estimates are calculated by first estimating the 2SLS regression, defined by equations (4.3) and (4.4), using unearned income as the endogeneous variable. Unearned income under the annuitization method is given by the annuity of the lottery win. Unearned income in the data is calculated as the difference between per-adult consumption and winner's earnings net of taxes and transfers. We then take cohort-size weighted averages of $\beta^{w,\ell}$ for each event time ℓ, and then taking the mean across estimates for event times 3-5. To calculate the intensive-margin MPE, we subtract the component that is due to responses along the extensive margin, using the extensive-margin shares reported in Table 4.2. We report intensive-margin MPEs for each subsample of winners with low-, medium- and high average incomes in the two years prior to winning.