Measuring the Insurance Value of Income Taxes

Sami Jysmä*

January 9, 2025

Abstract

Progressive taxes redistribute income from the rich to the poor, but also provide income insurance by redistributing income from periods of high income to periods of low income within the life-cycle of an individual taxpayer. This paper provides a framework for characterizing and measuring this insurance value. I provide a novel Slutsky-style decomposition of individual-level welfare impacts of tax changes when there is income uncertainty. Using this decomposition, I characterize optimal taxes with both redistributive preferences and uncertainty and the MVPF of tax reforms under income uncertainty. The willingness-to-pay (WTP) for the insurance aspect of taxes is a function of unobservable marginal utilities. Building on the unemployment insurance literature in the Baily (1978) – Chetty (2006) -tradition, I develop three methods to estimate the WTP. Using one of these, a novel consumption based approximation, I estimate the ex-ante MVPF of marginal tax increases at different points of the income distribution for the United States. The results indicate that when the insurance value is taken into account, the efficiency losses of tax increases are lower than otherwise, and those losses are decreasing rather than increasing with income.

Keywords: income uncertainty, redistribution, social welfare, MVPF, optimal tax, optimal social insurance

JEL Codes: H21, H23, H24, H55, I30

*Labor Institute for Economic Research LABORE and Tampere University. sami.jysma@labore.fi. I thank Raj Chetty, Jarkko Harju, Vesa-Matti Heikkuri, Nathan Hendren, Ohto Kanninen, Kaisa Kotakorpi, Arash Nekoei, Jukka Pirttilä, Stefanie Stantcheva and participants at seminars at Labore, Helsinki GSE, and Tampere University for their valuable comments.

1 Introduction

Progressive income taxes redistribute income from big earners to those with little income. At the same time, they also redistribute income from periods of high income to periods of low income within the life-cycle of an individual taxpayer. Taking into account the insurance value this provides is crucial for assessing the full welfare impacts of taxes. This has been acknowledged in the theoretical optimal tax literature for half a century (e.g. Mirrlees (1974), Varian (1980), Eaton and Rosen (1980)). Nevertheless, there are only a handful of papers trying to estimate the insurance value of income taxes (Hoynes and Luttmer (2011) and Stepner (2019)). This is in a stark contrast with the vast literature estimating the elasticity of taxable income (ETI), which provides a measure of the fiscal externalities of income taxation (see e.g. Neisser (2021) for a recent survey and references therein) or the insurance value of unemployment insurance programs (e.g. Gruber (1997), Chetty (2008), Hendren (2017), Landais and Spinnewijn (2021)).

A likely reason for the lack of empirical evidence on the insurance value of income taxes is a lack of theoretical work able to provide concrete and estimable characterizations of the insurance value. As a result, the theoretical and empirical literatures on the subject are mostly separate. This is again in contrast to the guidance provided by the theoretical literature on the importance of ETI and the insurance value of social insurance, which provide clear mappings from empirically estimable parameters to theoretical results.

This paper provides a framework linking results from welfare theory to empirically estimable characterizations of the insurance value of income taxes. I first provide a novel Slutsky-style decomposition of individual welfare effects of tax changes when there is income uncertainty. The decomposition consists of two parts, the *lifetime income effect*, and the *insurance effect*. With multiple tax payers, these effects have intuitive interpretations: the life-time income effects redistributes income between individuals, while the insurance effect redistributes income within the lifetime of an individual. Moreover, at the individual level, the effects are analogous to the income and substitution effects in a standard Slutsky decomposition of behavioral responses to price changes. The lifetime income effect simply characterizes how much the tax change affects an individual's lifetime income in the absence of behavioral effects. This is the only individual-level welfare impact in the absence of income uncertainty (e.g. in a static Mirrlees-model). The insurance effect then characterizes individual's willingness-to-pay (WTP) for the insurance aspect of a tax change that reduces the variation of post-tax income.

The decomposition allows me to express the welfare impacts of tax changes under income uncertainty in terms of individual WTP. Hence, the marginal value of public funds (MVPF) framework can be used by simply adding the total WTP for the insurance effect to the standard formula (e.g. Hendren (2016) and Hendren and Sprung-Keyser (2020)). Moreover, the decomposition makes it relatively simple to augment results from the static optimal tax literature to incorporate the insurance value of taxes: this only requires aggregating individual welfare impacts using social marginal welfare weights. I do this in the case of optimal non-linear income taxes in the lines of Diamond (1998) and Saez (2001). The decomposition to lifetime income and insurance effects proves valuable here: the redistribution of income between individuals through the lifetime income effects plays an analogous role to the redistribution of income in static models. The insurance effects then affect the optimal tax schedule through two conceptually different channels. The insurance value that taxes provide impacts the tax schedule through an efficiency rationale. However, it also has a redistributive rationale: there are differences in how much each tax payer benefits from the insurance provided so while the insurance effects do not redistribute income between individuals, they do redistribute welfare between individuals. The impact of this redistribution of welfare on the optimal tax schedule is then dependent on the welfare weights of the individuals.

Similarly to the Baily (1978) - Chetty (2006) literature, the insurance value of taxes is a function of unobserved marginal utilities. Here, the insurance value of a marginal tax increase at income level z, denoted by $\beta(z)$, is the percentage difference in marginal utilities below versus above that point times the probability of earning less than z. Drawing from the unemployment insurance (UI) literature (E.g. Baily (1978), Gruber (1997), Chetty (2006), Chetty (2008), Hendren (2017), Landais and Spinnewijn (2021)), I develop three ways of estimating $\beta(z)$. The first approach provides a novel consumption-based approximation that has attractive properties compared to the more traditional approximation used in the Baily (1978) - Chetty (2006) literature[1]: while the theory suggests that the insurance value is bounded between -1 and 1 as long as the marginal utility of consumption is higher below an income level z than above it – as it is with the novel approximation – the traditional approximation is unbounded. Hence, using the traditional approximation could lead to widely biased estimates, and it would do so in my empirical application.[2] For the second approach, I build on the work of Landais and Spinnewijn (2021) to show that $\beta(z)$ can be bounded from below if we have estimates of the marginal propensity to consume (MPC) at different points in the income distribution. The third approach further develops this result: unlike in Landais and Spinnewijn (2021), who consider responses only at the extensive margin, the continuity of earned income responses in my setting allows me to express the same lower bound in terms of the marginal propensity to earn (MPE) at different points in the income distribution. This is important as earnings data is much more widely available than consumption data.

Finally, I apply the methodology to estimate the ex-ante MVPF of marginal tax increases at different income levels in the United States. Ex-ante here refers to a situation where no information on individuals has yet been revealed to them.[3] I use the novel consumption-based approach together with public-use data for the year 2019 from the distributional national accounts of Piketty et al. (2018). Taking the insurance value of income taxes into account changes the ex-ante estimated MVPF

[1] Hendren (2021) shows that a a similar parameter can be approximated in the same way as has been done in the unemployment insurance literature.

[2] A similar approximation could also be applied in the Baily (1978) - Chetty (2006) framework.

[3] The methodology of this paper can be applied to any level of information revealed, but I choose to estimate the ex-ante MVPF due to the simplicity of the approach. See Hendren (2021) for a recent discussion of this point. Notably, maximizing ex-ante welfare is equivalent to maximizing ex-post utilitarian welfare.

substantially. Without the insurance value, the MVPF is above one at all income levels, suggesting that people are willing to pay more to avoid the tax increase than what the government obtains as tax revenue. In contrast with the insurance value taken into account, the MVPF is *lower* than one at all income levels even with low levels of relative risk-aversion. This means that people are willing to pay *less* to avoid marginal tax increases than what government receives from them. Moreover, without the insurance value MVPF estimates increase with income as the government receives less revenue from marginal tax increases towards the top. In contrast, with the insurance value, the MVPF is *decreasing* in income, so that willingness to pay for the tax hikes increases faster than tax revenue decreases. This suggests that marginal tax increases at the top would result in smaller efficiency losses than marginal tax increases at lower levels of income in an ex-ante sense.

Previous literature This paper relates to several strands of literature. Closest in terms of the research question is a rather scarce literature aiming to quantify the insurance value of taxes. To the best of my knowledge, there are two such papers: Hoynes and Luttmer (2011) and Stepner (2019). These papers, however, do not provide a link to widely used theoretical frameworks of optimal taxation and welfare analysis. Hence, the main contribution of this paper is to develop methods to estimate the insurance value of income taxes in a way that speaks directly to both widely used results in the optimal tax literature (esp. **?** and Saez (2001)), the optimal social insurance literature (esp. Baily (1978) and Chetty (2006)), and the MVPF literature (e.g. Hendren (2016), Hendren and Sprung-Keyser (2020)).

On the other side of the theory–empirics divide on this subject, this study relates to three strands of literature that deal with optimal unemployment insurance or taxation under uncertainty: the Baily (1978) – Chetty (2006) approach to optimal unemployment insurance, the Mirrlees (1974) – Varian (1980) approach to optimal taxation under uncertainty, and the more recent New Dynamic Public Finance (NDPF) literature on optimal dynamic taxation (e.g. Farhi and Werning (2013) and

Golosov et al. (2016)). The current paper differs from these strands of literature in several respects.

The most important difference to previous theoretical studies is that, as in the Baily (1978)–Chetty (2006) literature, the goal of this paper is to provide guidance for the empirical literature on how to estimate the insurance value of taxes. Previous studies in either the Mirrlees (1974)–Varian (1980) or the NDPF traditions do not express their results in terms of estimable sufficient statistics, and thus their impact on the empirical public finance literature has been modest. In contrast, this paper provides a direct link between theory and empirical work in this setting.

Another important difference is also that the setups analyzed in this paper differ in multiple ways from the setups analyzed in all of the previous strands. First, this paper analyzes multi-period life-cycle models where taxes are dependent only on the current labor income of each period. This is similar in spirit to the Baily (1978)–Chetty (2006) setup studying optimal social insurance which studies two state insurance. This paper then adds heterogeneous agents and redistributive preferences to that setup and also analyzes continuous labor income with general non-linear taxes. This setup is in contrast with both the Mirrlees (1974)–Varian (1980) literature, which analyzes two period models, and the NDPF literature which considers more general tax structures. In the sense that the tax systems analyzed in this paper are simple, this paper is similar to Findeisen and Sachs (2017), who analyze a life-cycle setting with a similar nonlinear labor tax as the one in chapter 3 of this paper, but with a linear tax on capital as well. While the goal of that paper is to provide theoretical results for optimal taxes, the goal here is to provide estimable sufficient statistics that relate directly to results in welfare theory.

Second, while the Baily (1978)–Chetty (2006) literature does not consider redistributive concerns in its canonical form, in both the Mirrlees (1974)–Varian (1980) and NDPF literatures the social planner's redistributive preferences are a function of only some of the individual shocks: in the Mirrlees (1974)–Varian (1980) literature redistributive preferences are defined over exogenous differences called "ability" and not over subsequent exogenous shocks, which are labeled "luck"[4]. In the NDPF literature, redistributive preferences depend only on the first period shock realization. In contrast, the current paper allows the social planner to have redistributive preferences that are a function of the full vector of life-time shocks faced by the individuals. To illustrate the importance of this difference, note that if the planner has redistributive tastes only over the first period shock, then those with the same first period shock have the same Pareto-weights regardless of any of their later shocks. Conditional on the first period shock, this constraints the social planner to consider only the efficiency rationale for shocks later in life. This clearly rules out many redistributive preferences that one might have.

In terms of methodology, this paper relates most closely to previous work on the perturbation approach to deriving optimal tax results in static models (e.g. Piketty (1997), Saez (2001), Saez (2002)), and on a similar approach to deriving optimal social insurance results following Baily (1978) and Chetty (2006). These literatures tend to address situations either in static models with heterogeneity (namely the static optimal tax models) or in dynamic models without heterogeneity (the Baily (1978) - Chetty (2006) approach). In contrast, I derive results here in situations with both heterogeneity and uncertainty. Indeed, the approach taken here can be seen as providing a unified way of analyzing many of these models.

The rest of the paper is organized as follows. Section 2 goes through a stylized model of unemployment insurance in order to give some intuition. Section 3 provides a more general treatment and considers nonlinear taxation. Section 4 discusses implementation. Section 5 provides the empirical application to ex-ante MVPF, and section 6 concludes.

[4]This somewhat artificial distinction was highlighted already by Eaton and Rosen (1980).

2 A stylized example with an employment tax

In this section, I introduce the main theoretical results of the paper in a stylized model where there is only an employment tax payed by the empolyed and an unemployment benefit. This setup is similar to the classic Baily (1978) - Chetty (2006) case but with redistributive preferences. The model can also be thought of as an extension of a Saez (2002) extensive margin model with two states to a case with uncertainty. Although the model is very stylized, the main intuition behind the results generalizes to more complex settings, such as the continuous non-linear tax case in section 3.

2.1 Setup

Individuals Consider a continuous time economy. There are two types θ of individuals in the economy: half are low-skilled L and the other half are high-skilled H. There are no income differences between the types conditional on employment status. Individuals can either be unemployed and receive a benefit $-T^U$ or be employed and receive z in wage income and pay taxes T^E, so that net income while employed is $z - T^E$. The types only differ in their probabilities of employment. Low-skilled individuals spend a fraction $h_L(s_L)$ of their life employed, whereas high-skilled workers spend a fraction $h_H(s_H)$ employed, where s_θ is the search effort exerted by type θ. The total fraction of employed is then $h(s_L, s_H) = \frac{1}{2}h_L(s_L) + \frac{1}{2}h_H(s_H)$. Both types have the same time-separable instantaneous quasi-linear utility function over consumption c and search effort s: $v(c, s) = u(c) - s$, where u is increasing and concave. There is no time discounting nor savings in the economy. Hence, unemployed workers consume the unemployment benefit $c^U = -T^U$ and employed workers consume their net-of-tax wages $c^E = z - T^E$ each instant. The workers maximize expected utility by choosing their search effort conditional on type and then learn their type, which stays the same.

Social planner The social planner does not observe the types of individuals, but knows the distribution of types in the economy. The planner, however, observes the incomes and employment statuses of individuals.

The planner has social preferences over different types, as captured by a social welfare function (SWF)[5] of the form

$$SWF = \frac{1}{2} G\left(E_H(v)\right) + \frac{1}{2} G\left(E_L(v)\right) \tag{1}$$

where the weights ($\frac{1}{2}$) come from the equal shares of both types of workers in the economy. G is a positive and weakly concave function over expected utilities that captures the redistributive preferences of the social planner. Note that if G is the identity function, then the social planner is simply maximizing the ex ante expected utility of individuals.[6] $E_\theta(v)$ is the expected (or life-time) utility of type θ. The planner has an exogenous revenue requirement R, and chooses the tax schedule (T_U, T_E) to maximize SWF conditional on a balanced budget $(1-h)T_U + hT_E = R$.

2.2 Tax reform analysis

I consider the social welfare impact of a small increase $d\tau$ of the taxes payed by the employed. The reform changes the disposable income (and consumption in this setting) of the employed from c_E to $c_E - d\tau$. This setup is illustrated in figure 1, where the solid line from c_U to c_E shows the budget line before the tax reform and the dashed line from c_U to $c_E - d\tau$ shows the post-reform budget line. The derivation here closely follows Saez (2002).

Impact on the government budget The impact of the perturbation on the government budget is the sum of the mechanical impact of the tax perturbation and the impact of any behavioral responses on tax revenue. These are summarized in figure 2.

[5] Similar results could be obtained by using generalized social welfare weights as primitives instead of a social welfare function as proposed by Saez & Stantcheva (2016).

[6] In other words, the social planner is maximizing the expected utility of individuals before they know their type. Under rational expectations, this amounts to maximizing the ex post utilitarian social welfare function.

The mechanical impact dM of the reform on the government budget is simply the share of employed h in the economy times the additional taxes they have to pay $d\tau$. Note that this is the case regardless of how h is determined. It does not matter, whether h is the share of the employed in a static Saez (2002) economy without uncertainty, or if h is the expected share of employed in a Baily-Chetty setup with uncertainty – in both cases, and with any mix of the two, the mechanical impact on the government budget can be expressed simply as

$$dM = hd\tau. \tag{2}$$

The behavioral impact dB in this setting comes from the movement of individuals from employment to unemployment due to the increased taxes payed while employed. In figure 2 this is illustrated by a gray horizontal arrow from employment to unemployment at the bottom of the figure. The behavioral impact is equal to the change in the share of unemployed dh due to the reform times the lost tax revenue due to the movement to unemployment $T_U - T_E$:

$$dB = dh(T_U - T_E). \tag{3}$$

Denoting by $\eta = \frac{dh}{d(c^E - c^U)} \frac{c^E - c^U}{h}$ the elasticity of the share of employed h with regard to the difference in consumption when employed versus when unemployed $c^E - c^U$, we can then express the behavioral effect as

$$dB = -\eta \frac{T^E - T^U}{c^E - c^U} hd\tau. \tag{4}$$

Note that also the behavioral effect can be expressed this way regardless of whether the change in the share of employed dh comes from a model with uncertainty or from a model with heterogeneity (or both).

Hence, the full impact of the reform on the government budget is

$$dM + dB = hd\tau - \eta \frac{T^E - T^U}{c^E - c^U} hd\tau. \tag{5}$$

Again, this expression is robust to the question whether the variation in the income distribution comes from consistent variation between individuals (heterogeneity) or from variation within individuals across time (uncertainty). Hence, all differences in the expression for the impact of the reform on social welfare will come from the welfare effects of the reform on individuals and from how those individual effects are aggregated to social welfare. I will next analyze these.

Impact on individual welfare – A Slutsky-style decomposition The reform impacts individual welfare in two ways: first, it changes the relative amount of taxes payed when unemployed versus when employed, and second, it affects the total amount of taxes an individual pays. The reform can be decomposed to to partial reforms accordingly: first, there is an *insurance reform* that changes the relative amount of taxes paid while employed versus while unemployed but does not have any mechanical impact on life-time income, and second, there is a *life-time income reform* that has no mechanical impact on the relative amounts of taxes paid and only has a mechanical impact on life-time income. These reforms can be used to decompose the welfare impact of the full reform for a given individual to an *insurance effect* and a *life-time income effect*. These are analogous to the substitution and income effects on demand in the Slutsky equation, but the income and life-time income effects are for welfare. These are illustrated in figure 3.

The insurance reform is such that it has no mechanical impact on the life-time income of a given type θ: it simply moves income from employment to unemployment. This give rise to the insurance effect. Hence, it takes $(1 - h_\theta)d\tau$ units of income away from the individual while employed and gives $h_\theta d\tau$ units back when unemployed[7], where h_θ is the fraction of time type θ spends employed. These income

[7]The mechanical impact of this reform is then $h_\theta \left((1 - h_\theta)d\tau\right) - (1 - h_\theta)\left(h_\theta d\tau\right) = 0$

changes are valued by the individual by the marginal utilities of consumption while employed $u_E^c(\theta)$ and while unemployed $u_U^c(\theta)$. Hence, type θ's willingness to pay for the insurance reform in terms of life-time income, WTP_θ^I, is given by

$$WTP_\theta^I = \beta_\theta h_\theta d\tau$$

$$= \underbrace{h_\theta \frac{(1-h_\theta)d\tau u_E^c(\theta)}{E_\theta(u^c)}}_{\text{WTP for income change when employed}} + \underbrace{(1-h_\theta)\frac{h_\theta d\tau u_U^c(\theta)}{E_\theta(u^c)}}_{\text{WTP for income change when unemployed}} \tag{6}$$

$$= h_\theta(1-h_\theta)\frac{u_U^c(\theta) - u_E^c(\theta)}{E_\theta(u^c)}d\tau,$$

where $\beta_\theta = (1-h_\theta)\frac{u_U^c(\theta)-u_E^c(\theta)}{E(u^c(\theta))}$ is the percentage difference in marginal utilities while unemployed versus while employed times the probability of being unemployed. The insurance reform is illustrated by a shift from the original solid black budget line (from c_U to c_E) to the blue dashed line from c_U^I to c_E^I in figure 3.

The life-time income reform then has a mechanical impact on life-time income only and does not redistribute income from employment to unemployment for the given individual. It gives rise to the life-time income effect on individual welfare. This reform simply takes $h_\theta d\tau$ units of income from type θ both when employed and when unemployed. Hence, the willingness to pay for the life-time income reform in terms of life-time income, WTP_θ^R, is simply

$$WTP_\theta^R = -h_\theta d\tau. \tag{7}$$

The life-time income reform is illustrated by a shift from the blue dashed line (from $c_U^I to c_E^I$) to the red dashed line from c_U to $c_E - d\tau$ in figure 3.

Taking these two reforms together, the willingness to pay for the full reform WTP_θ

is simply

$$\begin{aligned} WTP_\theta &= WTP_\theta^I + WTP_\theta^R \\ &= \underbrace{\beta_\theta h_\theta d\tau}_{\text{Insurance effect}} \quad \underbrace{-h_\theta d\tau}_{\text{Life-time income effect}} . \end{aligned} \tag{8}$$

In other words, the full willingness to pay for the reform is the sum of the WTP for the insurance reform and the WTP for the life-time income reform.

Marginal value of public funds The marginal value of public funds (MVPF) is defined as the ratio of aggregate willingess-to-pay (WTP) for a given reform over its total cost to the government (e.g. Hendren (2016), Hendren and Sprung-Keyser (2020)). In this case that ratio is given by

$$MVPF = \frac{\text{Total willingness to pay}}{\text{Cost to government}} = \frac{\frac{1}{2}WTP_H + \frac{1}{2}WTP_L}{-(dM + dB)}$$

or

$$MVPF = \frac{1-\beta}{1-\eta\frac{T^E-T^U}{c^E-c^U}},$$

where $\beta = \frac{\frac{1}{2}h_L d\tau}{h d\tau}\beta_L + \frac{\frac{1}{2}h_H d\tau}{h d\tau}\beta_H$ is the weighted average of β_θ weighted by the mechanical impact of the reform on different types ($h_\theta d\tau$). Hence, compared to the static case, β is the only addition to this formula. This means that relative to the static model, tax increases have a lower MVPF as long as the marginal utility of income while unemployed is higher than the marginal utility of income while employed.[8]

From individual to social welfare The expression in equation 8 is at the individual level. Hence, we still need to aggregate this to the social welfare level to get the full welfare impact dW of the policy.

The social planner values individual welfare increases by the social marginal wel-

[8]And we are not at the wrong side of the Laffer-curve, so that the denominator is positive.

fare weight g_θ of the individuals. In the current setting, with the social welfare function defined in equation 1, g_θ is equal to $G'(E_\theta(v))E_\theta(u^c)/\lambda$, where λ is the shadow cost of government revenue. However, we could use generalized social marginal welfare weights (Saez and Stantcheva (2016)) as g_θ.

The social planner then sums over types to arrive at its valuation of all individual welfare impacts.

$$\begin{aligned} dW =& \frac{1}{2} g_H WTP_H + \frac{1}{2} g_L WTP_L \\ =& \frac{1}{2}(g_H \beta_H h_H d\tau - g_H h_H d\tau) \\ &+ \frac{1}{2}(g_L \beta_L h_L d\tau - g_L h_L d\tau) \\ =& \beta g^I h d\tau - g^R h d\tau. \end{aligned} \tag{9}$$

Here, β is the average of β_θ weighted by the mechanical impact of the reform on different types ($h_\theta d\tau$), g^I is the average of g_θ weighted by the willingness to pay for the insurance reform $\beta_\theta h_\theta d\tau$, and g^R is the average of g_θ weighted by the willingness to pay for the life-time income reform $-h_\theta d\tau$. g^R is analogous to the g_i term in Saez (2002). Note that β is positive as long as individual marginal utilities of income are higher while unemployed then while employed and there are some individuals that spend time in both states. It is also bounded above by 1. Also, note that while g^R tends to be between 0 an 1 as the high ability type gets more weight (i.e. they spend more time employed), g^I is above 1 if the low ability type has a larger insurance effect.

To further separate the insurance and redistributive motives of the planner, I express equation 9 as

$$dW = \underbrace{-g^R h d\tau}_{\text{Redistribution}} \underbrace{+\beta h d\tau}_{\text{Insurance}} + \underbrace{\beta(g^I - 1) h d\tau}_{\text{Interaction}}. \tag{10}$$

This separates the social welfare impact of the reform dW to three parts: 1) a redistribution term $-g^R h d\tau$ that depends solely on redistributive preferences, 2) an insurance term $\beta h d\tau$ that depends solely on the insurance value, and 3) an interaction term

$\beta(g^I - 1)hd\tau$ that depends on the presence of both.

Optimal taxes Throughout the paper, I will characterize optimal taxes using ABC-type formulas as opposed to providing Baily-Chetty-type conditions, as ABC-formulas are applicable to a larger range of situations. However, I provide a discussion of the relationship between these two approaches in Appendix E.

The full impact of the reform on social welfare is then given by the sum of the impact on the government budget ($dM + dB$) and the individual welfare impacts aggregated to social welfare (dW).

$$\begin{aligned} dSWF &= dM + dB + dW \\ &= hd\tau - \eta \frac{T^E - T^U}{c^E - c^U} hd\tau - g^R hd\tau + \beta hd\tau + \beta(g^I - 1)hd\tau \end{aligned} \tag{11}$$

Setting this to zero, we can characterize the optimum by

$$\frac{T^E - T^U}{c^E - c^U} = \underbrace{\frac{1}{\eta}}_{\text{Fiscal externality}} \left[\underbrace{1 - g^R}_{\text{Redistribution}} \underbrace{+\beta}_{\text{Insurance}} + \underbrace{\beta(g^I - 1)}_{\text{Interaction}} \right]. \tag{12}$$

This expression resembles that of the extensive-margin case in Saez (2002), but with some added terms in the square brackets. First, $1-g^R$ determines how the redistribution of life-time income between individuals due to the reform affects the optimum in isolation. This is analogous to the $1 - g_i$ term in Saez (2002). Second, β determines how insurance considerations impact the optimum in isolation from redistributive concerns. Finally, $\beta(g^I - 1)$ determines how these two rationals interact with each other: if those who benefit from the insurance aspect of the reform are also those who the government values more than average (i.e. $g^I > 1$), this term is positive and the employment tax is higher than without this interaction. Note that if the policy does not have any insurance value on average, so that β is zero, we have the same expression as in Saez (2002) and the optimum is based solely on the trade-off between redistribution (summarized by g^R) and efficiency (summarized by η). If, on the other hand, the planner has no redistributive preferences (i.e. $g_\theta = 1$ for all θ, so that $g^R = g^I = 1$), the optimum is solely based on the trade-off between insurance and efficiency. In Appendix E, I show that in this two state case without redistributive preferences, equation 12 together with the government budget constraint are then equivalent to the classic Baily–Chetty formula. However, when we have both insurance value and redistributive preferences, these two interact to change the optimum if insurance values and social marginal welfare weights are correlated.

3 Welfare effects of non-linear taxes

This section applies the toolkit introduced in the previous example to the problem of welfare analysis of non-linear taxes in a life-cycle model. I derive the welfare impacts of non-linear taxes using perturbation methods as in e.g. Saez (2001).

3.1 Setup

Individuals Individuals face a shock variable θ_t in periods $t = 1, ..., T$ that follows an arbitrary process. The shock can be multidimensional and captures any aspects that impact the individual's utility, such information on past shocks or expectations of future shocks, changes in tastes etc. I define θ as a particular life-time shock an individual might face (i.e. $\theta = (\theta_1, ..., \theta_T)$). I will refer to θ as the type of an individual. Note that in period t an individual of type θ observes only the t first entries of θ (i.e. $\theta_1, ..., \theta_t$). In each period individuals observe the shock realization of that period and then make their choices.

Utility in period t, $u_t(c_t, l_t, \theta_t)$[9], is a function of consumption in that period c_t, a vector of work effort variables l_t, and the shock term θ_t. For simplicity, I assume that there are no income effects. This will only change the expression for the behavioral impact of taxes on the government budget and importantly will not affect the expres-

[9]Note that the utility function here allows for discounting as we might have e.g. $u_t = \delta^t u$.

sion for the insurance value of taxes.[10] Individuals maximize their period-level utility subject to a budget constraint, $c_t = z(l_t, \theta_t) - T(z(l_t, \theta_t)) + m(\theta_t)$, where $z_t = z(l_t, \theta_t)$ is pre-tax income, $T(z_t)$ is a (non-linear) tax function, and $m_t = m(\theta_t)$ is unearned income. In the model in the main text, I assume that the budget constraint in a given period is independent from the work effort and consumption of other periods, so that there are no savings.

Social planner The social planner knows the distribution of types, $F(\theta)$ in the economy, but does not observe the types of individuals. However, the planner observes the pre-tax income z_t of each individual. By choosing a nonlinear tax schedule $T(z)$, the planner maximizes a social welfare function of the form

$$SWF = \int_\theta G\left(\sum_t u_t(c_t, l_t, \theta_t)\right) dF(\theta) \tag{13}$$

subject to individuals' maximization and an external revenue requirement on the government budget $\int_\theta \sum_t T(z_{\theta_t}) dF(\theta) = R$. I assume that the planner chooses the tax schedule before any information has been revealed and does not change it afterwards. In other words, the tax schedule is assumed to be time-independent. Moreover, the tax schedule is assumed to be a function of an individual's current period earnings only[11] The function G is a positive and weakly concave function that characterizes the redistributive preferences of the planner. Again, if G is the identity function, the social planner is simply maximizing the ex ante expected utility of individuals before any information has been revealed. Hence, the functional form in 13 allows for ex-ante preferences for redistribution based on the type θ of individuals.

[10] Income effects could be added as in Saez (2001).

[11] This is in contrast with the new dynamic public finance literature (e.g. Golosov et al. (2016)), where the tax schedule is allowed to be a function of past earnings among other observables. Hence, the tax system here is constrained to be relatively simple, although this is a relatively close approximation of actual tax systems.

3.2 Reform and impact on the government budget

Reform Consider a similar tax reform as in Saez (2001). That is, I consider a tax reform within a small interval $[z, z+dz]$, such that marginal taxes within that interval are increased by $d\tau$. Hence, in periods when individuals earn more than $z + dz$, they pay $d\tau dz$ more taxes. This is illustrated in figure 4. The reform has two types of impacts: first, it changes the amount of taxes the government collects, so that it has an impact on the government budget, and second, it changes individual welfare. I will next derive expressions for these impacts.

Impact on the government budget The impact on the government budget can be expressed in the same way as in the static case (i.e. Saez (2001)). To see this, note that from the point of view of government revenue, it does not matter whether some individuals earn above z only some of the time or whether they earn some given $z^* > z$ all of the time. What matters for the government budget is how much there are taxpayers earning above z, and thus paying the extra $dzd\tau$ in taxes, and how individuals change their behavior due to the reform. The impact on the government budget is illustrated in figure 5.

The mechanical impact of the tax reform in this case is simply the mass of individual-periods spent earning more than z times the additional taxes paid due to the reform (in the absence of any behavioral responses). This means that the mechanical effect is simply

$$\begin{aligned} dM &= \int_\theta \sum_t d\tau dz \mathbb{1}(z_{\theta_t} > z) dF(\theta) \\ &= \int_\theta \sum_t \mathbb{1}(z_{\theta_t} > z) dF(\theta) d\tau dz \\ &= \int_\theta [1 - H_\theta(z)] dF(\theta) d\tau dz \\ &= [1 - H(z)] d\tau dz, \end{aligned} \tag{14}$$

where $\mathbb{1}(z_{\theta_t} > z)$ is an indicator function of whether individual of type θ would earn above z without the tax reform in period t, $H_\theta(z)$ is the CDF of period-level earnings of type θ at z, and $H(z) = \int_\theta H_\theta(z)dF(\theta)$ is the CDF of all period-level earnings in the economy at z.

As we have assumed no income effects, the reform will impact an individual's behavior only when their income would be within the interval $[z, z + dz]$ without the reform (as illustrated by a grey arrow to the left within that interval in figure 5). Next, denote the change of pre-tax income of type θ in period t due to the reform by $dz_{\theta_t}(z)$. In periods when the pre-tax income of type θ is within $[z, z + dz]$, this can be then expressed as

$$dz_{\theta_t}(z) = \varepsilon_{\theta_t}(z)z\frac{d\tau + dT'(z)}{1 - T'(z)} = \varepsilon_{\theta_t}(z)z\frac{d\tau + T''(z)dz_{\theta_t}(z)}{1 - T'(z)},$$

where $\varepsilon_{\theta_t}(z) = \frac{dz_{\theta_t}(z)}{z}/\frac{d(1-T'(z))}{1-T'(z)}$ is the period t elasticity of taxable income of type θ at income level z with regard to the marginal net-of-tax rate.[12] Solving for $dz_{\theta_t}(z)$ and using the virtual density $h^*_\theta(z)$ (as in Saez (2001)), we have that the behavioral effect for type θ is given by

$$dB_\theta = -\varepsilon_\theta(z)z\frac{T'(z)}{1 - T'(z)}h^*_\theta(z)d\tau dz,$$

where $\varepsilon_\theta(z) = \sum_t \varepsilon_{\theta_t}(z)\frac{\mathbb{1}(z_{\theta_t}=z)}{h^*_\theta(z)}$ is the average elasticity of taxable income of type θ at earned income z.

The full behavioral response is then given by the total response across types

$$\begin{aligned} dB &= \int_\theta dB_\theta dF(\theta) \\ &= -\varepsilon(z)\frac{T'(z)}{1 - T'(z)}\alpha(z)[1 - H(z)]d\tau dz, \end{aligned} \tag{15}$$

where $\varepsilon(z) = \int_\theta \varepsilon_\theta(z)\frac{h^*_\theta(z)}{\int_\theta h^*_\theta(z)}$ is the average elasticity of taxable income at z and $\alpha(z) = z\int_\theta h^*_\theta(z)/[1 - H(z)]$ is the local Pareto-parameter at z.

[12] Note that in periods where type θ earns more or less income than z, $dz_{\theta_t}(z)$ is 0 and hence also $\varepsilon_{\theta_t}(z) = 0$.

3.3 Slutsky-style decomposition of individual welfare effects

Due to uncertainty in the form of the shock term θ_t, the welfare effect of the reform differs from the static case. Next, I decompose the welfare impact for a given type θ to two separate effects: the *insurance effect* and the *life-time income effect*. As in the stylized model of section 2, these arise as the welfare impacts of two type-dependent partial reforms: the insurance reform and the life-time income reform, which are illustrated in figure 6. The insurance reform (in blue in the figure) changes the marginal tax rates in the same way as the full reform, but compensates the individual for any mechanical impact on the individuals life-time income. Hence, as an individual of type θ would spend $1-H_\theta(z)$ of their life earning more than z without the reform, the insurance reform would change the marginal tax rates to those in the full reform, but lower the tax schedule by $[1-H_\theta(z)]d\tau dz$ at each point. The life-time income reform (in red in the figure) on the other hand would not change marginal tax rates, but would increase taxes by the mechanical effect of the full tax reform on the life-time income of the individual, so that it would increase taxes by $[1-H_\theta(z)]d\tau dz$. Together these reforms amount to the full reform.

Insurance effect The insurance reform takes $H_\theta(z)d\tau dz$ units of income from an individual of type θ when they would have earned above z without the reform and gives $[1-H_\theta(z)]d\tau dz$ units of income back when they would have earned less than z without the reform. This ensures that the reform has no mechanical impact on the life-time income of the individual, as they receive $[1-H_\theta(z)][H_\theta(z)d\tau dz]$ units and are taken $H_\theta(z)[[1-H_\theta(z)]d\tau dz]$ unis from. The individual's willingess-to-pay for the insurance reform in terms of life-time income is then

$$
\begin{aligned}
&WTP^I_\theta(z) = \\
&\underbrace{H_\theta(z)\frac{[1-H_\theta(z)]d\tau dz E_\theta(u^c_{\theta_t}|z_{\theta_t}\le z)}{E_\theta(u^c_{\theta_t})}}_{\text{WTP below } z} - \underbrace{[1-H_\theta(z)]\frac{H_\theta(z)d\tau dz E_\theta(u^c_{\theta_t}|z_{\theta_t}> z)}{E_\theta(u^c_{\theta_t})}}_{\text{WTP above } z} \qquad (16)\\
&= [1-H_\theta(z)]d\tau dz\beta_\theta(z),
\end{aligned}
$$

where $E_\theta(u^c_{\theta_t}|z_{\theta_t} \leq z)$, $E_\theta(u^c_{\theta_t}|z_{\theta_t} \leq z)$ and $E_\theta(u^c_{\theta_t})$ are the expected marginal utilities of income of type θ below z, above z and across the whole income distribution of the type. The term $\beta_\theta(z) = H_\theta(z)\frac{E_\theta(u^c_{\theta_t}|z_{\theta_t}>z)-E_\theta(u^c_{\theta_t}|z_{\theta_t}\leq z)}{E_\theta(u^c_{\theta_t})}$ is the share of life spent earning less than z, $(H_\theta(z))$, times the percentage difference in expected marginal utilities of income below versus above z.

Life-time income effect The life-time income effect of type θ in terms of willingness-to-pay is simply the mechanical amount of additional taxes the individual has to pay due to the tax reform. Hence, this is simply given by

$$WTP^R_\theta(z) = -[1 - H_\theta(z)]d\tau dz. \tag{17}$$

Full individual welfare impact The full individual level willingness-to-pay for the reform is then simply the sum of the WTP for the partial reforms, so that

$$\begin{aligned} WTP_\theta(z) &= WTP^I_\theta(z) + WTP^R_\theta(z) \\ &= \underbrace{[1 - H_\theta(z)]d\tau dz\beta_\theta(z)}_{\text{Insurance effect}} \quad \underbrace{-[1 - H_\theta(z)]d\tau dz}_{\text{life-time income effect}}. \end{aligned} \tag{18}$$

3.4 From individual to social welfare

To go from individual willingness-to-pay for the reform to the social welfare impact dW, we must aggregate the individual WTP's based on how much the social planner values a dollar given to each type compared to that dollar being equally distributed across everyone. This is captured by the social marginal welfare weights of the planner, which are given by $g_\theta = \frac{G'(E_\theta(u_{\theta_t}))E_\theta(u^c_{\theta_t})}{\lambda}$, where λ is the shadow price of government revenue. Hence, the full social welfare effect of the reform is

$$\begin{aligned}
dW &= \int_\theta g_\theta (WTP^I_\theta(z) + WTP^R_\theta(z))dF(\theta) \\
&= WTP^I(z) \int_\theta g_\theta \frac{WTP^I_\theta(z)}{WTP^I(z)} dF(\theta) + WTP^R(z) \int_\theta g_\theta \frac{WTP^R_\theta(z)}{WTP^R(z)} dF(\theta) \\
&= \beta(z)[1-H(z)]d\tau dz g^I(z) - [1-H(z)]d\tau dz g^R(z) \\
&= \left(\beta(z) - g^R(z) + \beta(g^I(z)-1)\right)[1-H(z)]d\tau dz.
\end{aligned} \tag{19}$$

Here, $WTP^I(z)$ and $WTP^R(z)$ are the aggregate WTP's for each partial reform. $\beta(z) = \int_\theta \beta_\theta(z) \frac{[1-H_\theta(z)]d\tau dz}{[1-H(z)]d\tau dz} dF(\theta)$ is the weighted average of $\beta_\theta(z)$ where the weights are the mechanical effects of the reform on the life-time income of each type (so that $\beta(z)[1-H(z)]d\tau dz$ is the aggregate WTP for the insurance reform). $g^I(z)$ and $g^R(z)$ are weighted averages of g_θ with weights $WTP^I_\theta(z)$ and $WTP^R_\theta(z)$.

The $g^R(z)$ term is similar to the $g(z)$ term in Saez (2001), so that it tells us how much the social planner values a dollar of life-time income given to individuals when they are earning above z compared to the value of that dollar when distributed equally across all individuals. The $g^I(z)$ term instead tells us how much the planner values a dollar given to those who benefit from the insurance aspect of the reform.

3.5 Marginal value of public funds

Given that we have both the impact on the government budget (i.e. $dM + dB$) and the aggregate willingness to pay for the tax reform (i.e. $WTP^I + WTP^R$), it is very simple to express the MVPF of the reform. The MVPF is defined (e.g. Hendren (2016), Hendren and Sprung-Keyser (2020)) to be

$$MVPF = \frac{\text{WTP}}{\text{Cost to government}}.$$

Thus, in this situation, we can express the MVPF of the reform at z as

$$MVPF(z) = \frac{1-\beta(z)}{1-\varepsilon(z)\frac{T'(z)}{1-T'(z)}\alpha(z)}, \tag{20}$$

where $\alpha(z) = \frac{zh^*(z)}{1-H(z)}$. Note that the only difference to the static case in this expression is the $\beta(z)$ term. Hence, compared to the static case, the MVPF of such a tax reform is lower as long as marginal utility of income is higher below z than above it and we are not on the wrong side of the Laffer-curve (so that the cost to the government from a tax increase is negative, or in other words that an increase in the tax rate increases government revenue).

3.6 Optimal non-linear taxes

Optimal taxes in this case can be characterized from the FOC of the planner. That is, at the optimum, we must have that $dM + dB + dW = 0$. Using expressions for these from equations 14, 15 and 19 and simplifying, the optimal tax schedule can be characterized by

$$\frac{T'(z)}{1-T'(z)} = \frac{1}{\varepsilon(z)\alpha(z)}\left[\underbrace{1-g^R(z)}_{\text{Redistribution}} + \underbrace{\beta(z)}_{\text{Insurance}} + \underbrace{\beta(z)(g^I(z)-1)}_{\text{Interaction}}\right]. \tag{21}$$

Note that this expression is very similar to the formula in the simple two state case in the section 2 (equation 12). Here also $g^R(z)$ is similar to the $g(z)$ term in Saez (2001). In other words, it characterizes how much the social planner values an additional dollar in the hands of those, who earn more than z in a given period compared to that dollar being evenly distributed among everyone. The pure insurance rationale for taxation is captured by $\beta(z)$. Moreover, there is an interaction term $\beta(z)(g^I(z)-1)$ that is larger than 0 if those who benefit from the insurance aspect of the reform are also such that the planner values their income more than the average (i.e. if $g^I(z) > 1$) (and negative if the planner values less).

4 Implementation

The main problem one faces when trying to empirically implement the MVPF in equation 20 and the optimal tax schedule in equation 21 is the estimation of $\beta(z)$, which is the percentage difference in marginal utilities above and below z times the probability of earning less than z. In this section, I provide three ways of estimating $\beta(z)$, building on the unemployment insurance literature where multiple methods have been developed to estimate the insurance value of two-state insurance (e.g. Gruber (1997), Chetty (2008), Hendren (2017), Landais and Spinnewijn (2021)).

The first two approaches are part of a family of so-called *optimization-based* approaches (see Landais and Spinnewijn (2021)). The first of these builds on Landais and Spinnewijn (2021) and provides a lower bound of the insurance value as a function of the marginal propensity to consume out of unearned income (MPC) at different points of the income distribution. The second method is more novel and shows that with continuous earnings responses, one can also estimate a lower bound as a function of the marginal propensity to *earn* out of unearned income (MPE) at different points of the income distribution. This is important, as earnings data is much more readily available then detailed consumption data.

Unfortunately, while the optimization-based approaches have attractive features, such as not having to rely on an assumption of stable preferences at different states, they are very demanding on the causal effect estimates used to approximate the insurance value. In contrast to a standard social insurance setup, where such estimates would have to be estimated separately for two states, estimating $\beta(z)$ across the income distribution requires estimates of MPC or MPE at *each point* of the income distribution. Due to the lack of such estimates, I here rely on a *consumption-based* approach for measuring the insurance value of income taxes that builds on the Baily–Chetty tradition[13]. In this approach, assuming that every one has the same utility function and that marginal utility of consumption is only dependent on the amount of consumption, one can approximate $\beta(z)$ in terms of consumption above and below

[13] See also Hendren (2021).

z with a relative risk-aversion parameter.

4.1 Optimization-based approaches to estimating lower bounds for $\beta(z)$

The basic setup here is the same as in section 3. In each period t an individual faces a shock variable θ_t, and then chooses how much effort l_t to exert and how much to consume c_t. Here, I will assume that the shock and the chosen effort are scalars that linearly determine how much the agent earns $z(l_t, \theta_t) = \theta_t l_t$. Utility in period t is $u_t(c_t, l_t, \theta_t)$ and the budget constraint is $c_t = z(l_t, \theta_t) - T(z(l_t, \theta_t)) + m(\theta_t)$, where $m(\theta_t)$ is unearned income.

In Appendix C I prove the following two propositions that provide ways to bound $\beta(z)$ below under similar assumptions as those in Landais and Spinnewijn (2021):

Proposition 1. *Under assumptions C1–C4,*

$$\beta(z) \geq H(z) \frac{E\left(\frac{MPC}{1-MPC}\middle| z_\theta < z\right) - E\left(\frac{MPC}{1-MPC}\middle| z_\theta \geq z\right)}{E\left(\frac{MPC}{1-MPC}\right)}.$$

This proposition means that we can provide a lower bound for $\beta(z)$ as long as we have estimates of the marginal propensity to consume out of unearned income (MPC) for the full range of incomes. However, as consumption and earnings are related through the individual budget constraint, proposition 2 states that we can also provide a lower bound, in terms of the marginal propensity to earn out of unearned income (MPE):

Proposition 2. *Under assumptions C1–C4,*

$$\beta(z) \geq H(z) \frac{E\left(\frac{1/(1-\tau_n)+MPE}{MPE}\middle| z_\theta < z\right) - E\left(\frac{1/(1-\tau_n)+MPE}{MPE}\middle| z_\theta \geq z\right)}{E\left(\frac{1/(1-\tau_n)+MPE}{MPE}\right)},$$

where $\tau_n(z_\theta)$ is the marginal tax rate at z_θ.

4.2 Consumption-based approach to estimating $\beta(z)$

However, due to difficulties in estimating the MPC and the MPE for the full range of incomes, I here rely on a consumption-based approach. The approach I develop here is different to those often applied in the UI literature and to that proposed in Hendren (2021), who shows that a similar insurance parameter can be approximated by the percentage difference in consumption times a relative risk-aversion parameter γ. The reason why I develop another method is that the standard approximation is highly biased at high levels of $H(z)$, i.e. at the top end of the income distribution. In my empirical setting for example, the standard approximation would suggest that increasing marginal taxes at the top would have an MVPF of infinity, while that is impossible in the model.[14]

To arrive at an approximation for $\beta(z)$, I make the following assumption, which is often made in the UI literature (and is made in Hendren (2021):

Assumption 1. *The marginal utility of consumption depends only on the level of an individual's consumption, so that there exists a function f such that*

$$\frac{\partial u(c,l,\theta)}{\partial c} = f(c).$$

In Appendix D, I prove the following proposition, which provides an approximation of $\beta(z)$:

Proposition 3. *Under assumption 1*

$$\beta(z) \approx H(z)\frac{E\left(1/c_\theta^{\gamma_\theta}|z_\theta < z\right) - E(1/c_\theta^{\gamma_\theta}|z_\theta \geq z)}{E(1/c_\theta^{\gamma_\theta})}.$$

This means that we need data on consumption at different income levels and an estimate of the coefficient of relative risk-aversion at different levels of income. I next apply this approximation to estimating the ex-ante MVPF of marginal tax increases at different points of the income distribution.

[14] This is because $\beta(z)$ is between -1 and 1 always, and between 0 and 1 as long as the marginal utility of consumption is decreasing in taxable income.

5 Ex-ante MVPF's of marginal tax increases

In this section, I apply the methods developed in this paper to estimate the ex-ante marginal value of public funds (MVPF) of marginal tax increases at different points of the US wage income distribution. "Ex-ante" in this case refers to a situation where no information on individual's types is yet known. The methodology of this paper can be applied at any point of information revelation, but I choose to examine the ex-ante MVPF due to the clarity of the approach (see e.g. Hendren (2021) for a recent discussion on ex-ante welfare).

I assume that the joint distribution of earnings and disposable income is stable across time. This implies that a cross-sectional joint distribution of earnings and disposable incomes (taken as a measure of consumption) is equal to the joint probability distribution of those variables for all individuals ex-ante.[15] Hence, under this assumption and using the consumption based approach in proposition 3, one only needs a cross section of incomes and consumption together with an estimate of risk-aversion to estimate $\beta(z)$ and thus the willingness-to-pay for a marginal tax increase.

In this case, the ex-ante MVPF of a marginal tax increase at income-level z is given by

$$MVPF(z)_{ex-ante} = \frac{1-\beta(z)}{1+FE(z)} = \frac{1-\beta(z)}{1-\alpha(z)\varepsilon(z)\frac{T'(z)}{1-T'(z)}}, \tag{22}$$

where $\beta(z)$ is the insurance value and $FE(z) = -\alpha(z)\varepsilon(z)T'(z)/[1-T'(z)]$ is the fiscal externality. As this expression notes, one also needs the current tax-schedule ($T'(z)/[1-T'(z)]$) and estimates of $\alpha(z)$ and the taxable income elasticity $\varepsilon(z)$ in order to estimate the local MVPF.

[15] To see this, consider the stylized model in section 2. Before one knows one's type, the probability of earning z and consuming $c_E = z - T^E$ is $\frac{1}{2}h_L + \frac{1}{2}h_H$, and the probability of earning 0 and consuming $c_u = -T^U$ is $\frac{1}{2}(1-h_L) + \frac{1}{2}(1-h_H)$. But these are equal to h and $1-h$ respectively, or the shares of employed and unemployed in the economy.

Data As my data source, I use the U.S. Distributional National Accounts public use files for 2019 Piketty et al. (2018). I use the wage-income variable as my income measure, and extended disposable income as my consumption variable. I construct the marginal tax schedule for 2019 by using the NBER TAXSIM.

5.1 Results

Insurance value I use the consumption based-approach to measuring $\beta(z)$ developed in section 4, so that I approximate $\beta(z)$ by the expression in proposition 3 for different constant relative risk aversion parameters γ:

$$\beta(z) \approx H(z)\frac{E\left(1/c_\theta^{\gamma_\theta}|z_\theta < z\right) - E(1/c_\theta^{\gamma_\theta}|z_\theta \geq z)}{E(1/c_\theta^{\gamma_\theta})}.$$

Estimates of $1-\beta(z)$ at different 5,000 \$ bins with different risk aversion parameters are illustrated in figure 7. Here, we can see that as long as individuals are risk-averse, so that relative risk-aversion γ is positive, $1-\beta(z)$ approaches 0, but the exact rate of convergence is dependent on the value of relative risk-aversion. If, however, individuals are risk-neutral, so that $\gamma = 0$, $1-\beta(z)$ is 1 at all levels of income. Note that the risk-neutral case provides the same results as a case with no uncertainty. As we will see, this difference has large implications for the MVPF of marginal tax increases when comparing a situation where we take uncertainty into account to a situation where we do not.

Fiscal externality As in a standard case without uncertainty, the fiscal externality of a marginal tax increase at income-level z, $FE(z)$, is $-\alpha(z)\varepsilon(z)T'(z)/[1-T'(z)]$. I estimate $\alpha(z) = zh(z)/[1-H(z)]$ from the empirical income distribution at different 5,000\$ income bins and smooth it as the moving average of three bins above 100,000\$ due to the large variation in the point estimates (these are illustrated in the appendix figure A1). I then assume different constant levels of the elasticity $\varepsilon(z)$ between 0.1 and 0.5 and use marginal taxes provided by the NBER TAXSIM. The results for

$1 + FE(z)$ are illustrated in figure 8. There, we can see that $1 + FE(z)$ is decreasing in income up to around 100,000 $ of income and is fairly stable above it. The size of the elasticity scales the estimates, so that higher elasticities lead to lower $1 + FE(z)$ values. These results are consistent with past literature (e.g. Saez (2001)), where $\alpha(z)$ has been found to be stable at large levels of income. This then results in the stability of $1 + FE(z)$ also.

MVPF Finally, the ex-ante MVPF of marginal tax increases at income-level z is the ratio $[1 - \beta(z)]/[1 + FE(z)]$. Estimates of this are shown in figure 9 for elasticity $\epsilon(z) = 0.5$.[16] Strikingly, while the MVPF is above one and increasing with either risk-averse individuals or with no uncertainty ($\gamma = 0$), the MVPF is below one and decreasing with moderate values of risk-aversion. This implies that the efficiency costs of increasing marginal taxes are much lower when we take insurance-value into account. Moreover, the efficiency costs are *decreasing* in income with insurance value, so that at high levels of income, the efficiency cost is close to 0, while the government still receives tax income. This is in stark contrast with the case of no uncertainty, when the efficiency cost of marginal tax increases is *increasing*, so that for a given amount of additional revenue collected, individuals are willing to pay more to avoid the tax increase at higher levels of income. Taken at face value, this suggests that even at low levels of risk-aversion (e.g. with $\gamma = 1$), taxing high income individuals would be a more attractive way to obtain additional revenue compared to taxing low-income individuals purely from an efficiency point of view.

6 Conclusion

This paper shows how one can apply a simple Slutsky-style decomposition of individual welfare to characterize the welfare impacts of tax changes. This allows me to express the MVPF of marginal income tax changes and optimal income taxes under

[16] MVPF estimates with different assumed elasticity values are shown in Appendix figures A2 and A3

income uncertainty in terms of sufficient statistics. Building on the unemployment insurance literature in the Baily (1978) – Chetty (2006) tradition, I also develop methods to estimate a central parameter that arises in the welfare analysis, $\beta(z)$, which is the percentage difference of marginal utilities below versus above z times the probability of earning below z. I also apply this newly developed methodology to estimating the ex-ante MVPF of marginal tax increases for the US. Taking the insurance value of income taxes into account results in widely different estimates compared to the MVPF estimates without the insurance value: with the insurance value, the cost of marginal tax increases are lower at all income levels, and the costs decrease with income. This means that in an ex-ante sense, taking the insurance value into account, marginal tax increases induce substantially less efficiency costs and the efficiency costs are *decreasing* and not increasing in income.

The methodology developed in this paper can be used to guide further empirical work into the insurance value of taxes – a crucial parameter for welfare analysis taxes that has received very little attention compared to, for example, the empirical work on the elasticity of taxable income or the insurance value of explicit social insurance programs. Moreover, the framework of this paper could be easily adjusted to answering related questions, such as whether age-dependent taxation would provide welfare improvements through the insurance aspect, or whether taxes should be based on annual income or approximately on life-time income.

References

Baily, M. (1978): "Some Aspects of Optimal Unemployment Insurance," *Journal of Public Economics*, 10, 379–402.

Chetty, R. (2006): "A General Formula for the Optimal Level of Social Insurance," *Journal of Public Economics*, 90, 1879–1901.

——— (2008): "Moral Hazard versus Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, 116, 1150–1196.

Diamond, P. (1998): "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates," *American Economic Review*, 88, 83–95.

Eaton, J. and H. Rosen (1980): "Optimal Redistributive Taxation and Uncertainty," *Quarterly Journal of Economics*, 95, 357–364.

Farhi, E. and I. Werning (2013): "Insurance and Taxation over the Life Cycle," *Review of Economic Studies*, 80, 596–635.

Findeisen, S. and D. Sachs (2017): "Redistribution and Insurance with Simple Tax Instruments," *Journal of Public Economics*, 146, 58–78.

Golosov, M., M. Troshkin, and A. Tsyvinski (2016): "Redistribution and Social Insurance," *American Economic Review*, 106, 359–386.

Gruber, J. (1997): "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, 87, 192–205.

Hendren, N. (2016): "The Policy Elasticity," *Tax Policy and the Economy*, 30, 51–89.

——— (2017): "Knowledge of Future Job Loss and Implications for Unemployment Insurance," *American Economic Review*, 107, 1778–1823.

——— (2021): "Measuring Ex Ante Welfare in Insurance Markets," *The Review of Economic Studies*, 88, 1193–1223.

Hendren, N. and B. Sprung-Keyser (2020): "A Unified Analysis of Government Policies," *Quarterly Journal of Economics*, 135, 1209–1318.

Hoynes, H. and E. Luttmer (2011): "The Insurance Value of State Tax-and-Transfer Programs," *Journal of Public Economics*, 95, 1466–1484.

Landais, C. and J. Spinnewijn (2021): "The Value of Unemployment Insurance," *Review of Economic Studies*, 88, 3041–3085.

Mirrlees, J. (1974): "Notes on Welfare Economics, Information and Uncertainty," in *Essays in Equilibrium Behavior under Uncertainty*, ed. by D. M. M. Balch and S. Wu, Amsterdam, the Netherlands: Elsevier, North Holland.

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

Piketty, T. (1997): "La Redistribution Fiscale face au Chômage," *Revue Francaise d'Economie*, 12, 157–201.

Piketty, T., E. Saez, and G. Zucman (2018): "Distributional National Account: Methods and Estimates for the United States," *Quarterly Journal of Economics*, 133, 553–609.

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

——— (2002): "Optimal Income Transfer Programs," *Quarterly Journal of Economics*, 117, 1039–1073.

Saez, E. and S. Stantcheva (2016): "Generalized Social Marginal Welfare Weights for Optimal Tax Theory," *American Economic Review*, 106, 24–45.

Stepner, M. (2019): "The Insurance Value of Redistributive Taxes and Transfers," Working Paper.

Varian, H. (1980): "Redistributive Taxation as Social Insurance," *Journal of Public Economics*, 14, 49–68.

Figures and Tables

Figures

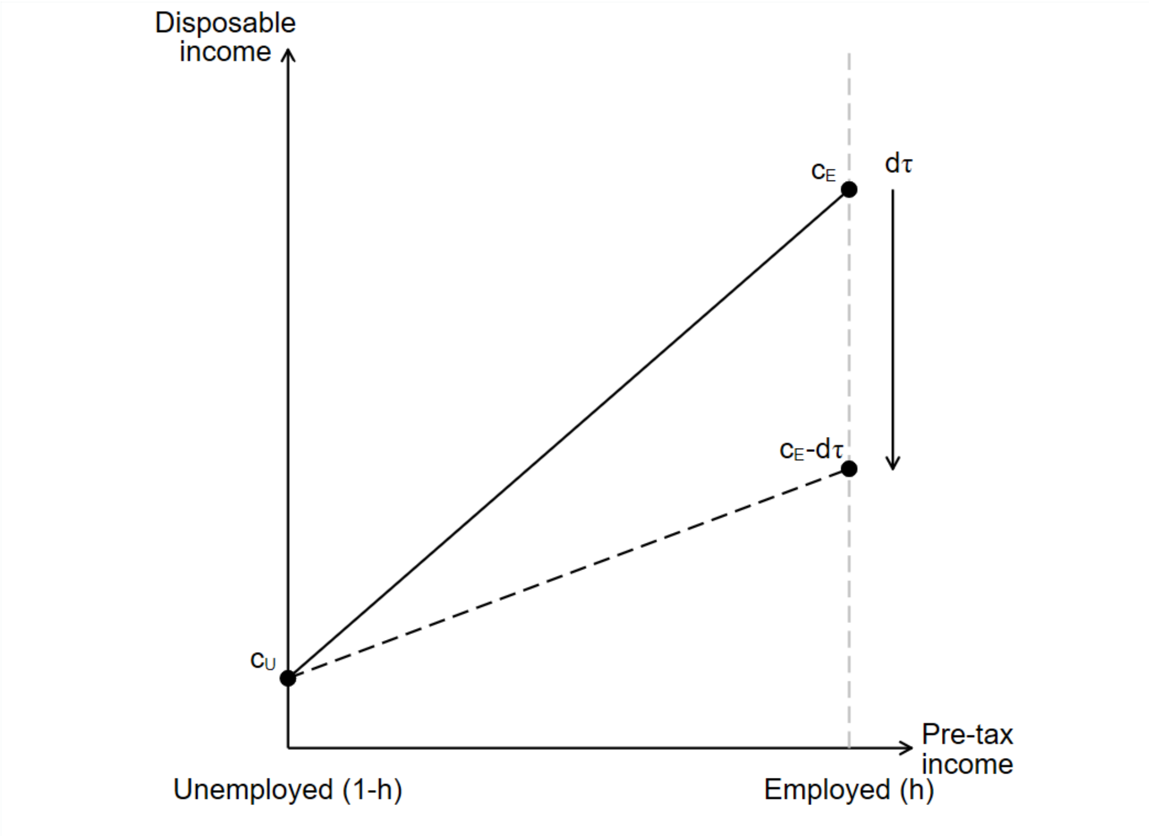

Figure 1: Tax reform $d\tau$

Notes: This figure illustrates the setting in a stylized model on unemployment insurance analyzed in section 2. The budget line shifts from the solid line (c_U, c_E) to the dashed line $(c_U, c_E - d\tau)$ when the taxes on the employed are increased by $d\tau$.

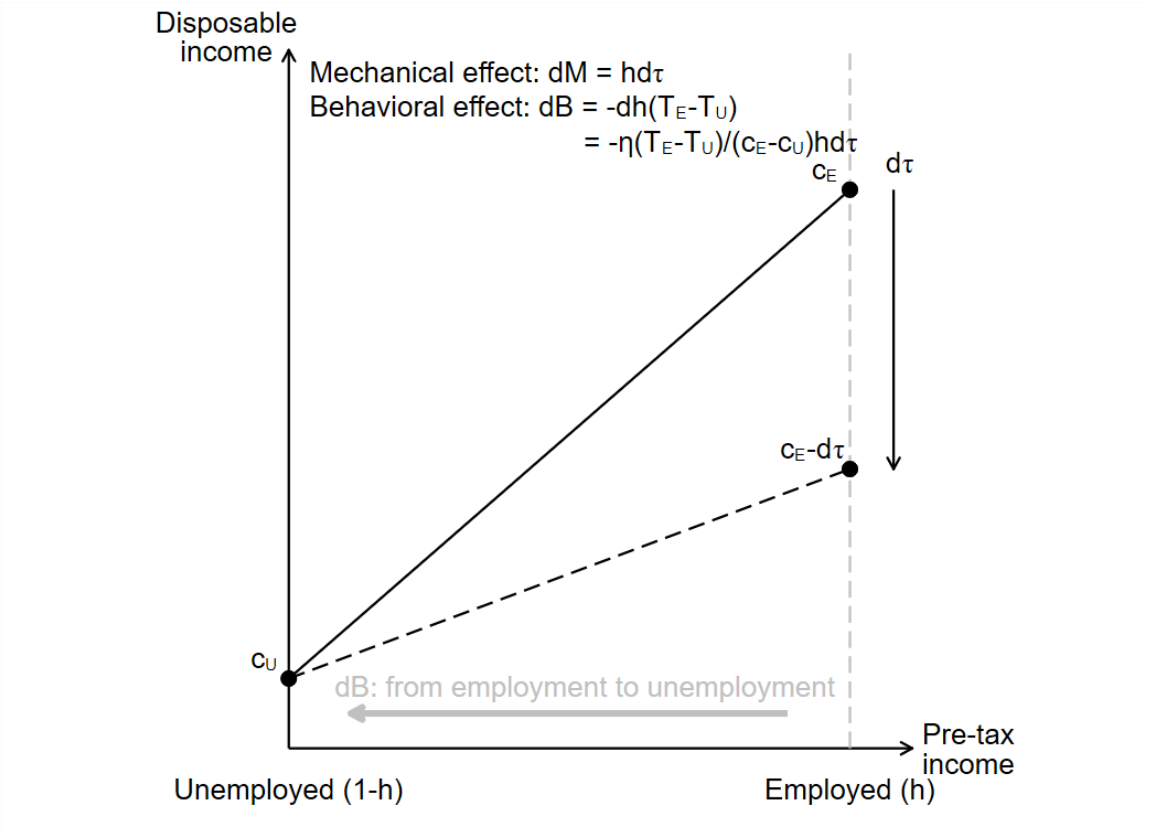

Figure 2: Impact of the tax reform on government revenue

Notes: This figure illustrates the impact of the increase in the taxes on the employed by $d\tau$ on the government budget. The mechanical effect is given by dM. Also, the tax change induces some to move from employment to unemployment as depicted by the grey horizontal arrow at the bottom. This induces a behavioral impact on the government budget dB.

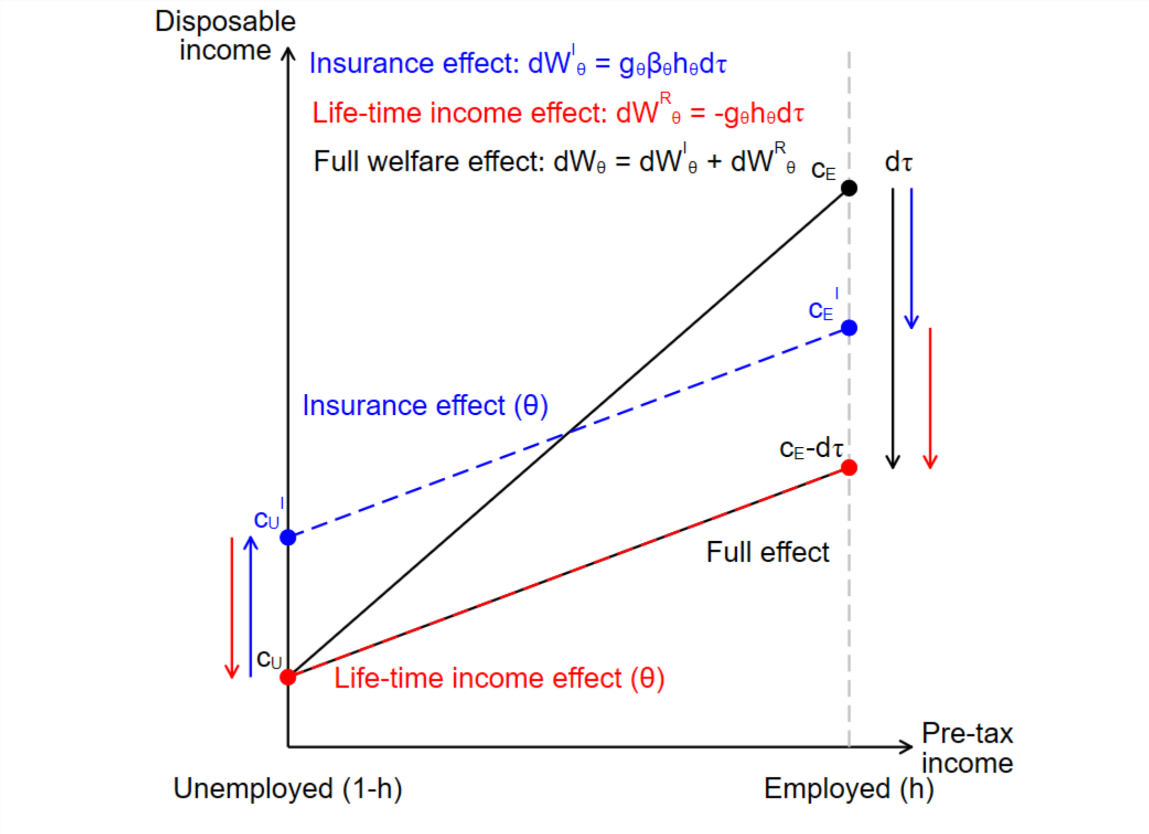

Figure 3: Impact on individual welfare

Notes: This figure illustrates the decomposition of the reform $d\tau$ at individual (θ) level to 1) an insurance reform (blue) and 2) a life-time income reform (red). The insurance reform redistributes income from times of employment to times of unemployment, but does not have a mechanical impact on the life-time income of the individual. Therefore, it pivots the budget line from the solid black line (c_U, c_E) to the blue dashed line (c_U^I, c_E^I). The life-time income reform then does not redistribute income from employment to unemployment, but only has a mechanical impact on life-time income and shifts the budget line from the blue dashed line (c_U^I, c_E^I) to the dashed red line $(c_U, c_E - d\tau)$. The individual welfare impact of the reform dW_θ equals the sum of the welfare impact of the insurance reform dW_θ^I and the life-time income reform dW_θ^R.

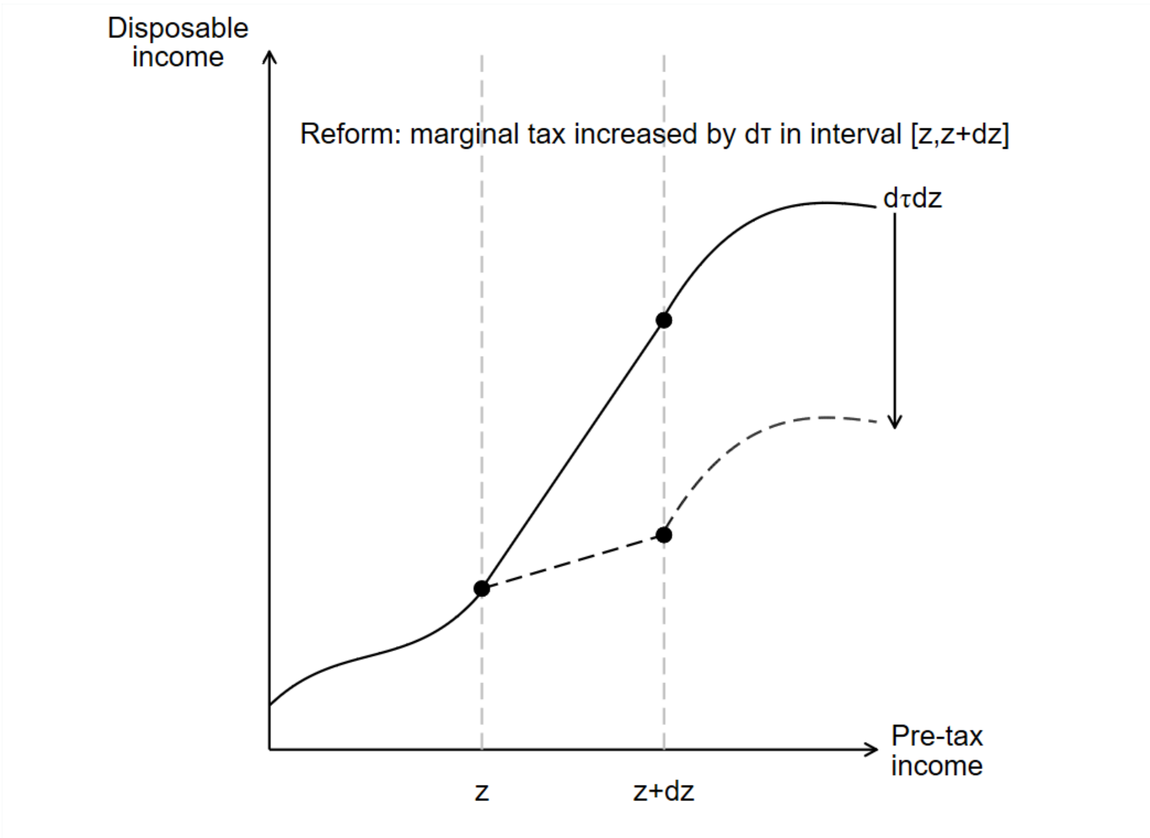

Figure 4: Nonlinear tax reform

Notes: This figure illustrates the tax reform (or tax perturbation) in the nonlinear case. The solid black line shows a continuous tax schedule before the reform. Then, in a small interval of pre-tax income from z to $z + dz$, marginal taxes are increased by $d\tau$. This means that disposable (or post-tax) income is decreased by $d\tau dz$ above $z + dz$, while the marginal tax rate changes only within the interval $[z, z + dz]$. The reform is similar to that in Saez (2001).

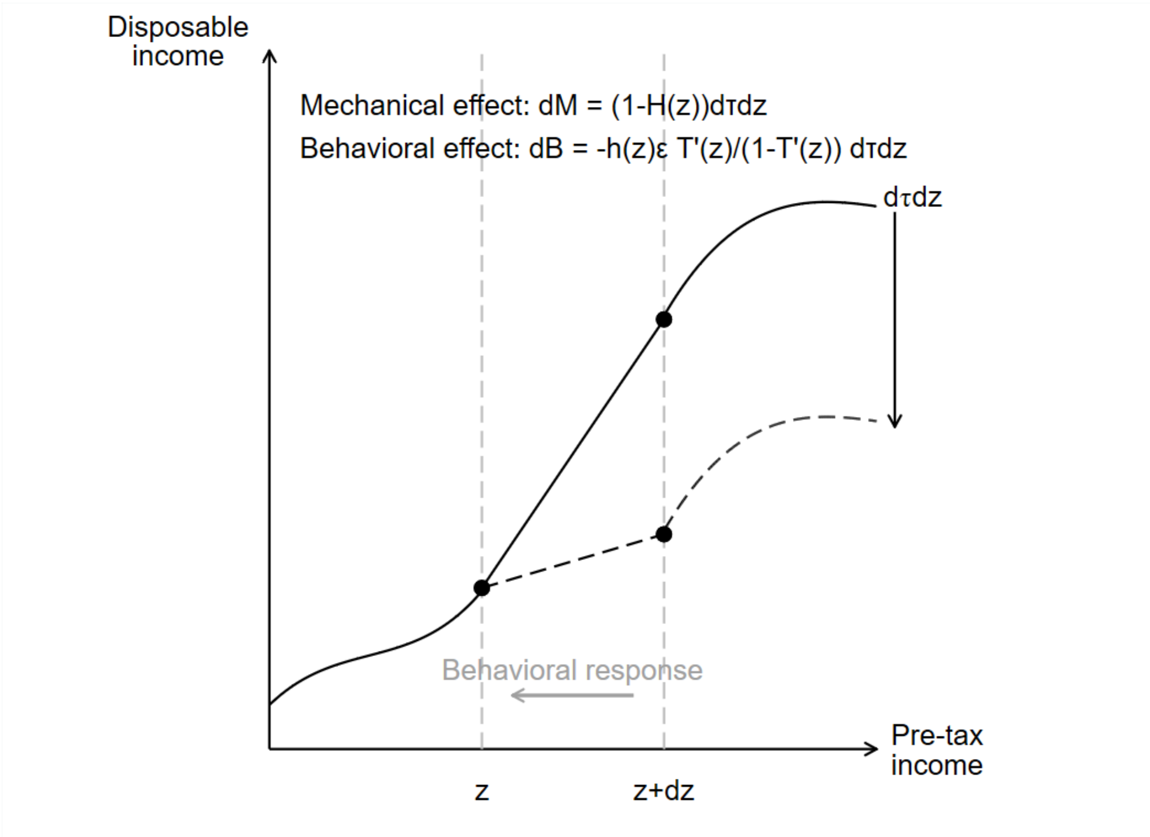

Figure 5: Nonlinear tax reform: impact on the government budget

Notes: This figure illustrates the impact of the tax reform of the govenrment budget. As in Saez (2001), there are two effects: the mechanical impact of the tax reform dM and the behavioral effect dB. The mechanical effect is simply the mass of time individuals spend above z during their whole life-time $1 - H(z)$ times the additional taxes they have to pay $d\tau dz$. As here we assume no income effects, individuals change their behavior only when they are within the interval $[z, z + dz]$ as this is the only interval where marginal taxes change. When individuals are in this interval, they reduce their taxable income, causing a fiscal externality on the government budget. These effects can be expressed in the same way as in Saez (2001).

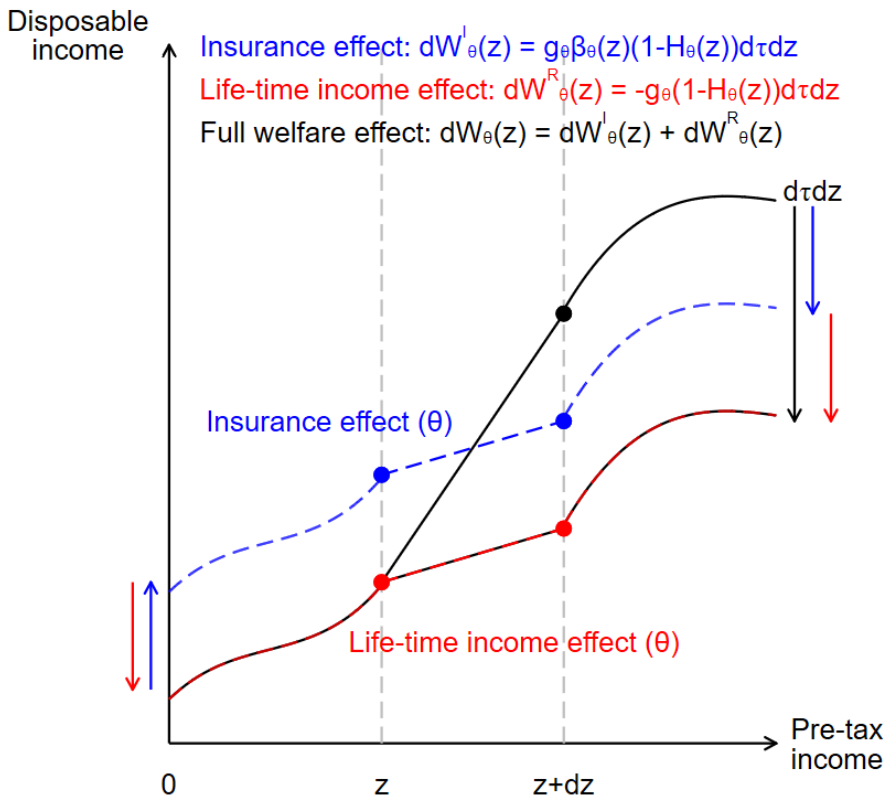

Figure 6: Nonlinear tax reform: Slutsky-style decomposition

Notes: This figure illustrates the Slutsky-style decomposition of individual welfare. For a given individual of type θ, the reform can be decomposed to two parts using two partial reforms: the insurance reform and the life-time income reform. For θ, the insurance reform changes the marginal tax rates in the same way as the full reform, but compensates for any mechanical impact on life-time income. The welfare effect of the insurance reform is referred to as the *insurance effect* as it depends solely on the insurance value that the full reform creates. The life-time income reform then does not change marginal tax rates, but adjusts the tax schedule as to have the same mechanical impact on the life-time income of θ as the full reform. The welfare effect of the life-time income reform is referred to as the *life-time income effect*. The full welfare impact of the reform on type θ is then simply the sum of the insurance and life-time income effects.

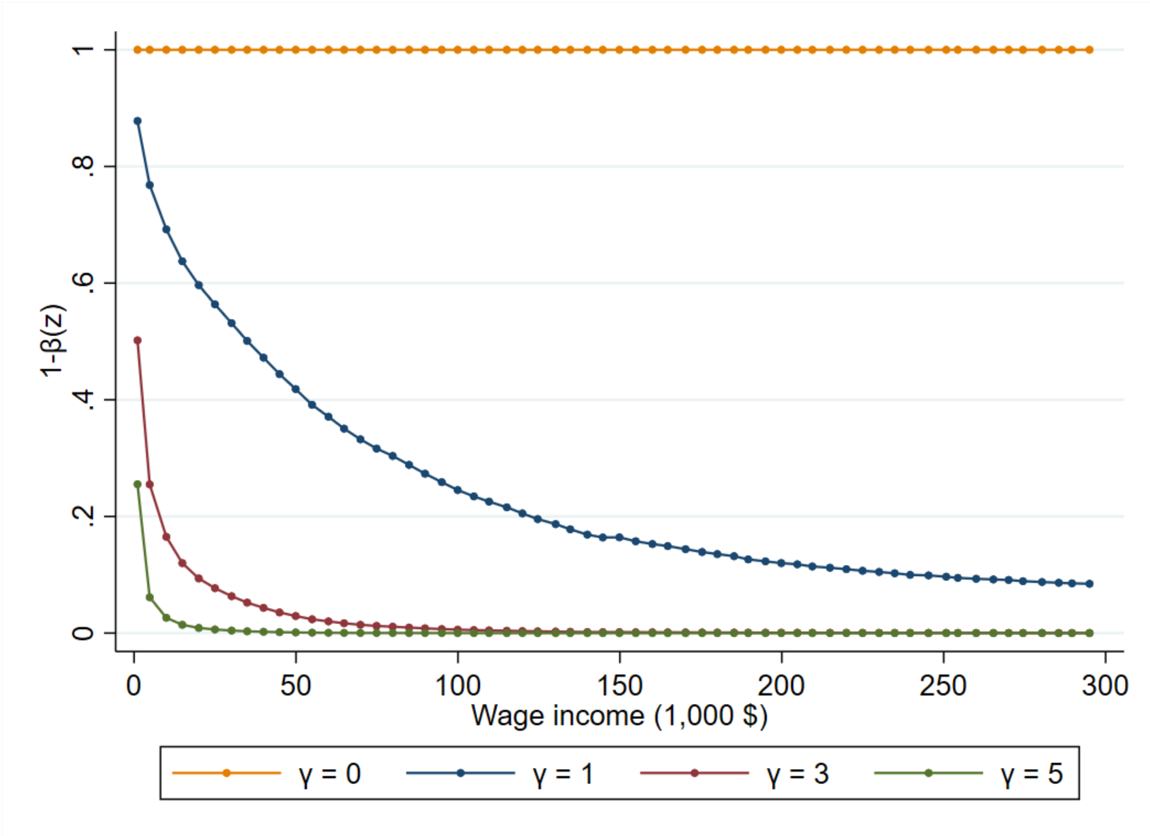

Figure 7: Estimates of the $1 - \beta(z)$

Notes: This figure shows estimates of $1 - \beta(z)$ for marginal tax-rate increases at 5,000$ bins from equation 22. $\beta(z)$ is the insurance value of increasing the marginal tax rate at income level z. It is the percentage difference of marginal utilities below and above z times the probability of earning below z. The figure shows estimates of $1-\beta(z)$ with different assumed levels of relative risk-aversion γ. With risk neutral individuals (i.e. when $\gamma = 0$), insurance value is 0 at all income levels, and $1-\beta(z)$ is simply 1. This corresponds to the case with no uncertainty. Instead, with positive values of risk-aversion, $1 - \beta(z)$ is less than 1 and converges to 0, so that the rate of convergence is faster with larger relative risk-aversion.

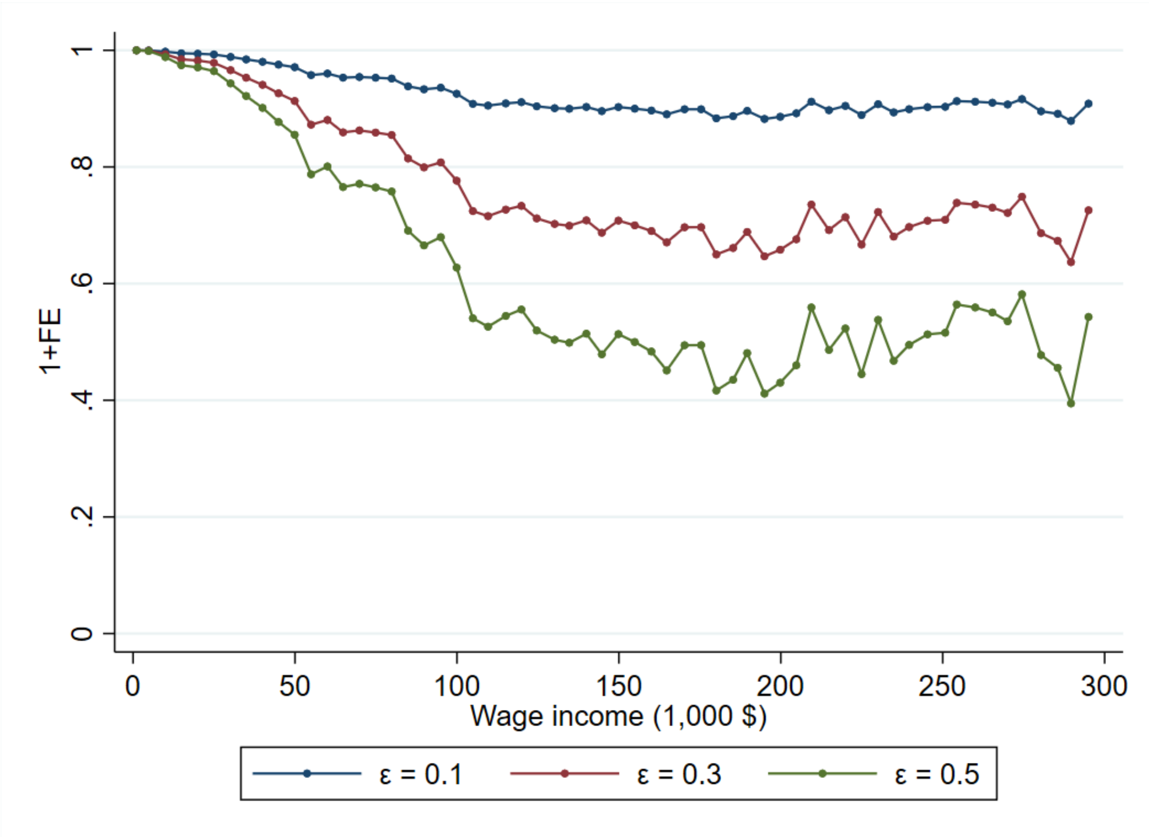

Figure 8: Estimates of $1 + FE(z)$

Notes: This figure shows estimates of $1 + FE(z)$ for marginal tax-rate increases at 5,000$ bins from equation 22. $FE(z) = \alpha(z)\varepsilon(z)T'(z)/[1 - T'(z)]$ is the fiscal externality of increasing the marginal tax rate at income level z. $\alpha(z)$ is a function of the income distribution, $\varepsilon(z)$ is the elasticity of taxable income, and $T'(z)$ is the marginal tax rate. The figure shows estimates of $1 + FE(z)$ with different assumed levels of the elasticity $\varepsilon(z)$.

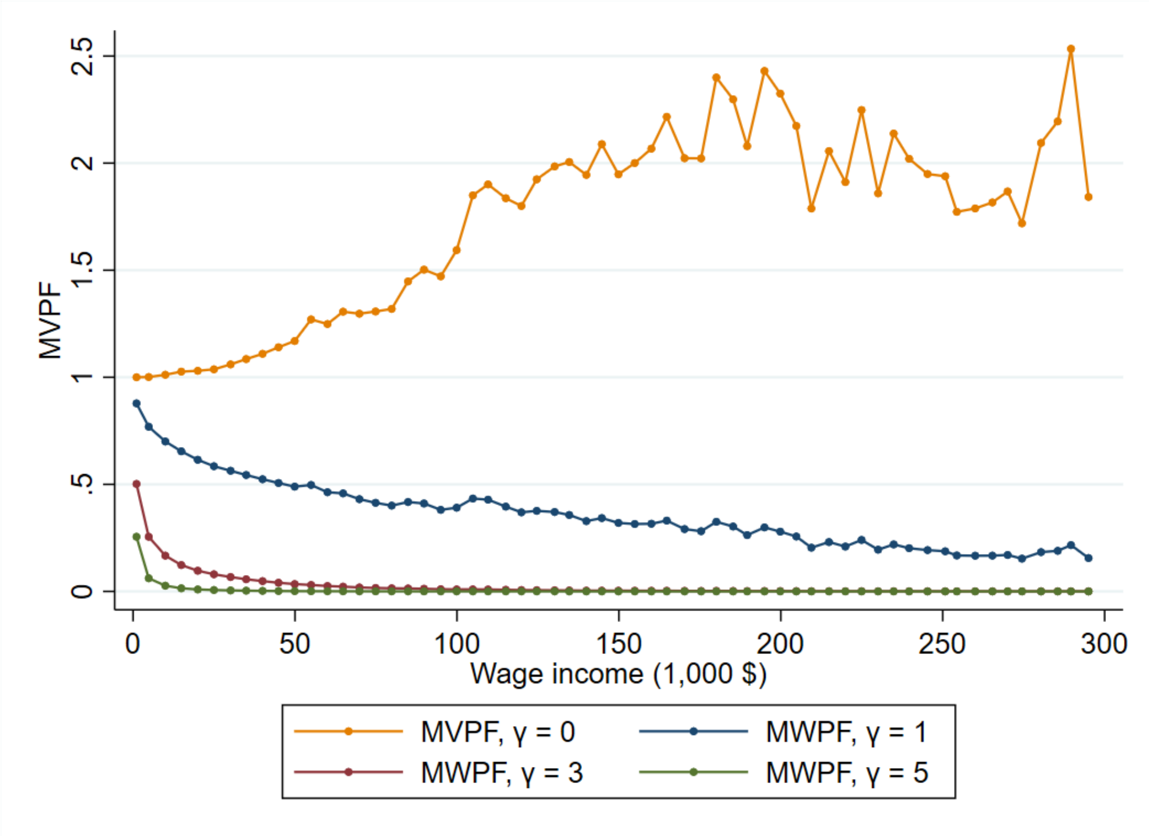

Figure 9: MVPF estimates with $\varepsilon(z) = 0.5$

Notes: This figure shows estimates of the ex-ante marginal value of public funds (MVPF) of marginal tax increases for different different assumed levels of relative risk-aversion γ and with an elasticity of taxable income $\varepsilon = 0.5$. $\gamma = 0$ corresponds to risk-neutral individuals, so that the MVPF is the same as for a situation without uncertainty. Moderate levels of risk-aversion substantially change the MVPF-estimates. Without the insurance value (i.e. with $\gamma = 0$), the MVPF is estimated to be *above* 1 at all income-levels, while for relative risk-aversion equal to 1, the MVPF is estimated to be *below* 1 at all income-levels. This implies that the efficiency costs of marginal tax increases are much lower when insurance value is taken into account. The MVPF estimates also switch from increasing with income to decreasing with income when insurance value is taken into account. This implies that when insurance value is taken into account, the efficiency costs of marginal tax increases would be lowest at the top, which is in stark contrast to the case without insurance value.

Appendices

A Appendix: Additional figures

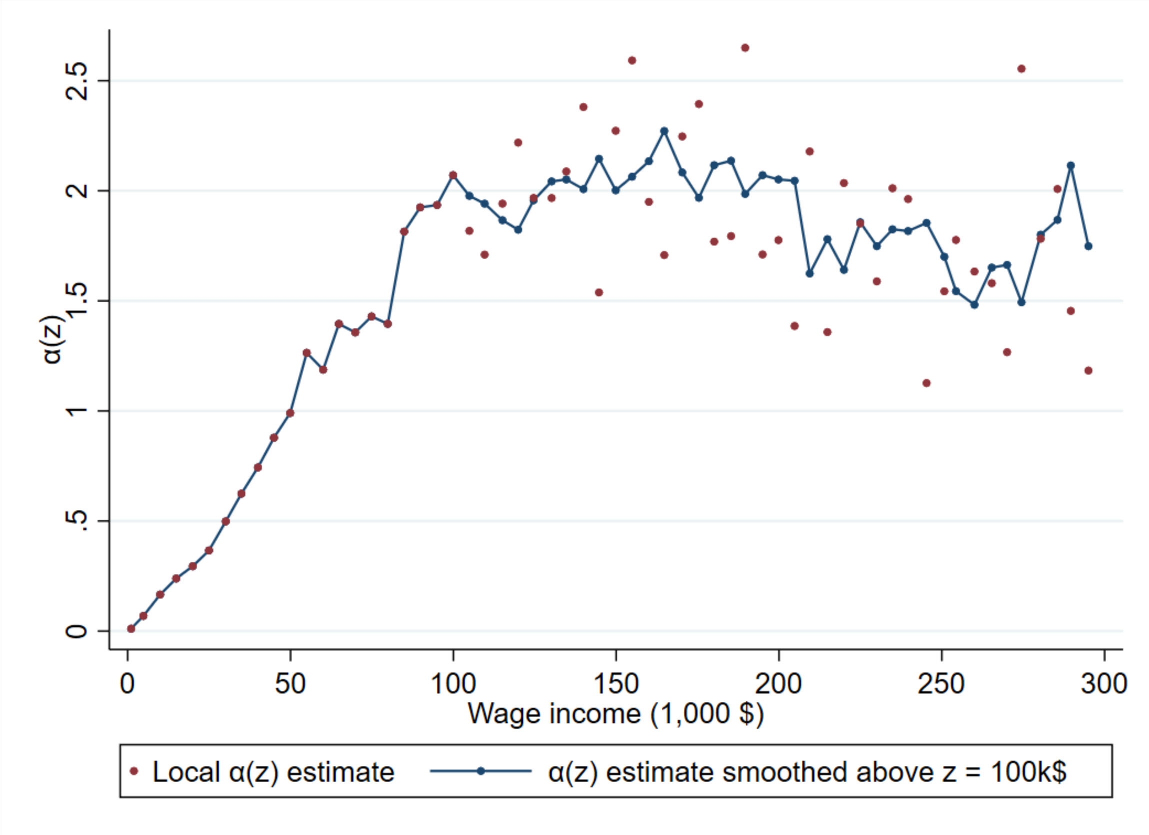

Figure A1: Estimates of $\alpha(z)$

Notes: This figure shows estimates for $\alpha(z) = h(z)/[1 - H(z)]$, where $h(z)$ is the probability distribution function and $H(z)$ is the cumulative distribution function of the income distribution. The red dots show local $\alpha(z)$ estimates without smoothing, and the blue line shows $\alpha(z)$ estimates that have been smoothed by a simple 3 dot moving average above 100,000 $ of income.

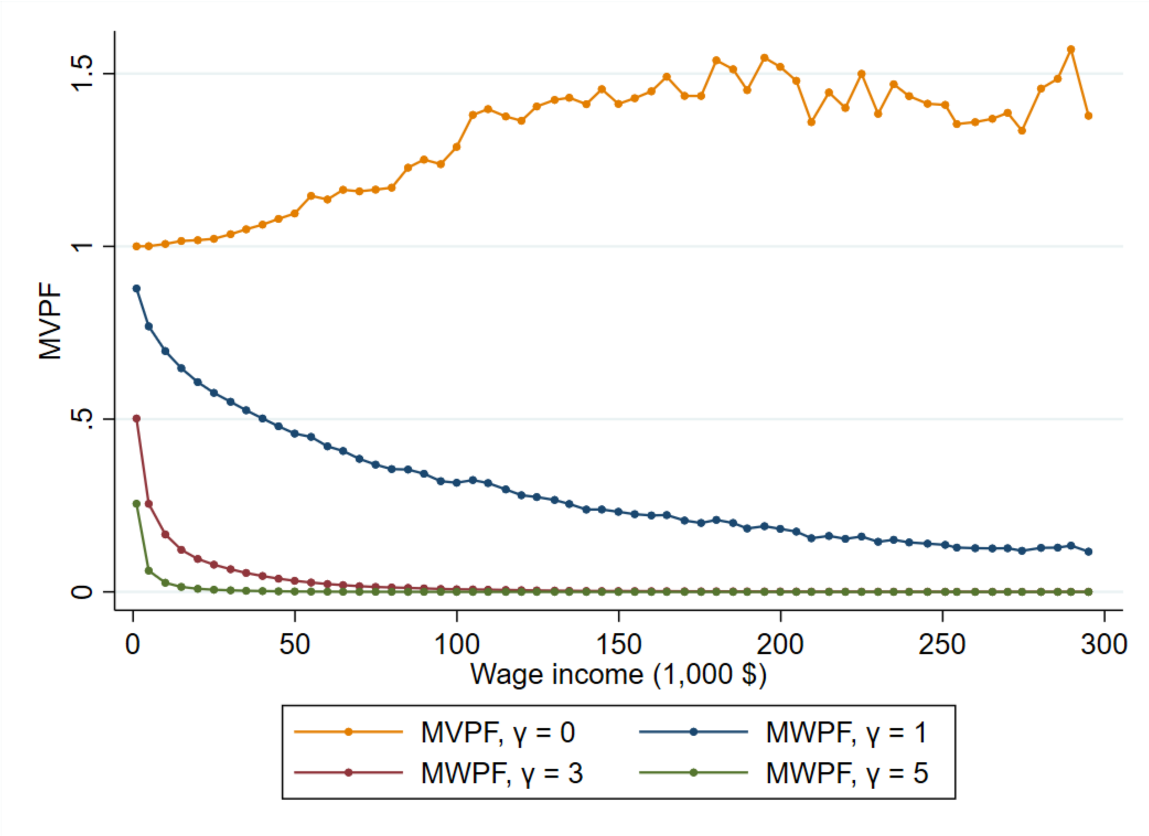

Figure A2: MVPF estimates with different levels of relative risk aversion and with $\varepsilon(z) = 0.3$

Notes: This figure shows estimates of the marginal value of public funds (MVPF) of marginal tax increases for different different assumed levels of relative risk-aversion γ and with an elasticity of taxable income $\varepsilon = 0.3$. $\gamma = 0$ corresponds to risk-neutral individuals, so that the MVPF is the same as for a situation without uncertainty. Moderate levels of risk-aversion substantially change the MVPF-estimates. Without the insurance value (i.e. with $\gamma = 0$), the MVPF is estimated to be *above* 1 at all income-levels, while for relative risk-aversion equal to 1, the MVPF is estimated to be *below* 1 at all income-levels. This implies that the efficiency costs of marginal tax increases are much lower when insurance value is taken into account. The MVPF estimates also switch from increasing with income to decreasing with income when insurance value is taken into account. This implies that when insurance value is taken into account, the efficiency costs of marginal tax increases would be lowest at the top, which is in stark contrast to the case without insurance value.

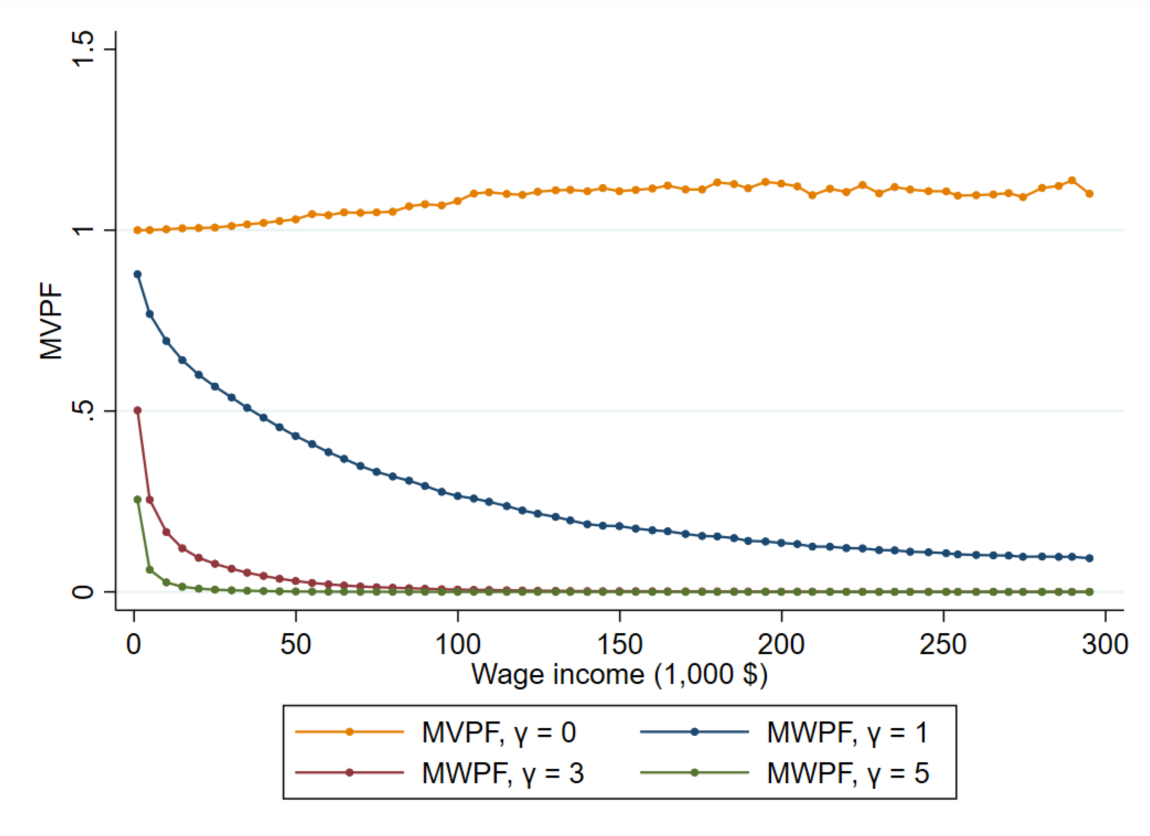

Figure A3: MVPF estimates with different levels of relative risk aversion and with $\varepsilon(z) = 0.1$

Notes: This figure shows estimates of the marginal value of public funds (MVPF) of marginal tax increases for different different assumed levels of relative risk-aversion γ and with an elasticity of taxable income $\varepsilon = 0.1$. $\gamma = 0$ corresponds to risk-neutral individuals, so that the MVPF is the same as for a situation without uncertainty. Moderate levels of risk-aversion substantially change the MVPF-estimates. Without the insurance value (i.e. with $\gamma = 0$), the MVPF is estimated to be *above* 1 at all income-levels, while for relative risk-aversion equal to 1, the MVPF is estimated to be *below* 1 at all income-levels. This implies that the efficiency costs of marginal tax increases are much lower when insurance value is taken into account. The MVPF estimates also switch from increasing with income to decreasing with income when insurance value is taken into account. This implies that when insurance value is taken into account, the efficiency costs of marginal tax increases would be lowest at the top, which is in stark contrast to the case without insurance value.

B Appendix: Formal derivation of equation 12

Let us study the problem of the social planner θ:

$$\max_{T_U, T_E} \int_\theta G\left((1-h_\theta)u(-T_U, e_\theta, \theta) + h_\theta u(z - T_E, e_\theta, \theta), \theta\right)$$
$$s.t. \int_\theta \left((1-h_\theta)T_U + h_\theta T_E\right) = 0$$

First order conditions (with envelope conditions taken into account):
For T_E:

$$\int_\theta \left[-G'_\theta h_\theta u^c_E(\theta) + \lambda\left(\frac{\partial h_\theta}{\partial T_E}(T_E - T_U) + h_\theta\right)\right] dT_E = 0 \tag{B1}$$

for T_U:

$$\int_\theta \left[-G'_\theta (1-h_\theta) u^c_U(\theta) + \lambda\left(\frac{\partial h_\theta}{\partial T_U}(T_E - T_U) + (1-h_\theta)\right)\right] dT_U = 0,$$

and the government budget constraint

$$\int_\theta \left((1-h_\theta)T_U + h_\theta T_E\right) = 0.$$

To arrive at the ABC-formula, we only need the FOC for T_E (or T_U). Dividing equation B1 by λ and dT_E, denoting $\eta_\theta = -\frac{\partial h_\theta}{\partial(c_E - c_U)}\frac{c_E - c_U}{h_\theta}$, and re-arranging we have

$$\int_\theta -\frac{G'_\theta E_\theta(u^c)}{\lambda}\frac{u^c_E(\theta)}{E_\theta(u^c)} h_\theta - \frac{T_E - T_U}{c_E - c_U}\int_\theta \eta_\theta h_\theta + \int_\theta h_\theta = 0.$$

Dividing by $h = \int_\theta h_\theta$, we get

$$\int_\theta -\frac{G'_\theta E_\theta(u^c)}{\lambda}\left(\frac{u^c_E(\theta)}{E_\theta(u^c)} - 1 + 1\right)\frac{h_\theta}{h} - \frac{T_E - T_U}{c_E - c_U}\int_\theta \eta_\theta \frac{h_\theta}{h} + 1 = 0.$$

$$
\begin{aligned}
\int_\theta -\frac{G'_\theta E_\theta(u^c)}{\lambda}\left(\frac{u^c_E(\theta)-h_\theta u^c_E(\theta)-(1+h_\theta)u^c_U(\theta)}{E_\theta(u^c)}\right)\frac{h_\theta}{h}& \\
+\int_\theta -\frac{G'_\theta E_\theta(u^c)}{\lambda}\frac{h_\theta}{h}-\frac{T_E-T_U}{c_E-c_U}\int_\theta \eta_\theta\frac{h_\theta}{h}+1=0.&
\end{aligned}
$$

$$
\begin{aligned}
\int_\theta \frac{G'_\theta E_\theta(u^c)}{\lambda}(1+h_\theta)\left(\frac{u^c_U(\theta)-u^c_E(\theta)}{E_\theta(u^c)}\right)\frac{h_\theta}{h}& \\
+\int_\theta -\frac{G'_\theta E_\theta(u^c)}{\lambda}\frac{h_\theta}{h}-\frac{T_E-T_U}{c_E-c_U}\int_\theta \eta_\theta\frac{h_\theta}{h}+1=0.&
\end{aligned}
$$

$$
\begin{aligned}
\frac{T_E-T_U}{c_E-c_U}\int_\theta \eta_\theta\frac{h_\theta}{h}=\int_\theta \frac{G'_\theta E_\theta(u^c)}{\lambda}(1+h_\theta)\left(\frac{u^c_U(\theta)-u^c_E(\theta)}{E_\theta(u^c)}\right)\frac{h_\theta}{h}& \\
1-\int_\theta \frac{G'_\theta E_\theta(u^c)}{\lambda}\frac{h_\theta}{h}.&
\end{aligned}
$$

$$
\begin{aligned}
\frac{T_E-T_U}{c_E-c_U}\int_\theta \eta=\beta\int_\theta \frac{G'_\theta E_\theta(u^c)}{\lambda}\frac{\beta_\theta}{\beta}\frac{h_\theta}{h}& \\
-\int_\theta \frac{G'_\theta E_\theta(u^c)}{\lambda}\frac{h_\theta}{h}+1.&
\end{aligned}
$$

$$
\frac{T_E-T_U}{c_E-c_U}=\frac{1}{\eta}\left(\beta g^I+1-g^R\right).
$$

$$
\frac{T_E-T_U}{c_E-c_U}=\frac{1}{\eta}\left(1-g^R+\beta-\beta(1-g^I)\right).
$$

C Appendix: Deriving propositions 1 and 2

The basic setup here is the same as in section 3. In each period t an individual faces a shock variable θ_t, and then chooses how much effort l_t to exert and how much to consume c_t. Here, I will assume that the shock and the chosen effort are scalars that linearly determine how much the agent earns $z(l_t, \theta_t) = \theta_t l_t$. Utility in period t is $u_t(c_t, l_t, \theta_t)$ and the budget constraint is $c_t = z(l_t, \theta_t) - T(z(l_t, \theta_t)) + m(\theta_t)$, where $m(\theta_t)$ is unearned income.

This gives the first order condition

$$\frac{\partial u_\theta}{\partial c} + \frac{1}{1 - T'(z_\theta)} \frac{1}{\theta} \frac{\partial u_\theta}{\partial l} = 0, \tag{C1}$$

where I have dropped the period t sub-index to save on notation and have emphasized that the utility-function and the earned income function vary by type θ.

I assume that $T(z)$ is piece-wise linear, so that as long as an agent is not located at a kink point $T'(z)$ is not affected by small changes in z.

Assumption C1. $T'(z)$ *is a step-function.*

This implies that there are N step-points $(z_1, ...z_n, z_{n+1}, ..., z_N)$, where the marginal tax rate changes. I denote the marginal tax rate between two adjacent step points (z_n, z_{n+1}) by τ_n. Hence, the budget constraint within the interval $[z_n, z_{n+1}]$ is

$$c_\theta = z_\theta - \tau_n[z_\theta - z_n] - T(z_n) + m_\theta, \tag{C2}$$

where $T(z_n)$ is the amount of taxes paid at z.

Then, if the optimum exists within the interval $]z_n, z_{n+1}[$, the first order condition C1 becomes

$$\frac{\partial u_\theta}{\partial c} + \frac{1}{1 - \tau_n} \frac{1}{\theta} \frac{\partial u_\theta}{\partial l} = 0. \tag{C3}$$

Consider next two shocks θ and θ' such that the individual chooses to earn less with shock θ. In other words, at the current optimum, $z_\theta \leq z_{\theta'}$. I assume that z_θ is between step-points z_n and z_{n+1}, and that $z_{\theta'}$ is between $z_{n'}$ and $z_{n'+1}$ such that $n' \leq n$. Using the optimality condition C3, we can express the marginal rate of substitution between these two states as

$$MRS_{\theta,\theta'} = \frac{\partial u_\theta/\partial c}{\partial u_{\theta'}/\partial c} = \frac{1-\tau_{n'}}{1-\tau_n}\frac{\theta'}{\theta}\frac{\partial u_\theta/\partial l}{\partial u_{\theta'}/\partial l}. \tag{C4}$$

As in Landais and Spinnewijn (2021), I assume that the marginal effort cost is higher in the worse state θ_t, so that

Assumption C2. $\frac{\partial u_\theta/\partial l}{\partial u_{\theta'}/\partial l} \geq 1.$

This assumption together with C4 readily implies that

$$MRS_{\theta,\theta'} \geq \frac{1-\tau_{n'}}{1-\tau_n}\frac{\theta'}{\theta}. \tag{C5}$$

I next derive estimable expressions for the RHS of equation C5. Taking a derivative of equation C3, with respect to unearned income m and using the budget constraint, we have

$$\begin{aligned}&\left[\frac{\partial^2 u_\theta}{\partial^2 c} + \frac{1}{1-\tau_n}\frac{1}{\theta}\frac{\partial^2 u_\theta}{\partial l \partial c}\right]\frac{\partial c_\theta}{\partial m}\\ &+\left[\frac{\partial^2 u_\theta}{\partial l \partial c} + \frac{1}{1-\tau_n}\frac{1}{\theta}\frac{\partial^2 u_\theta}{\partial^2 l}\right]\underbrace{\frac{1}{1-\tau_n}\frac{1}{\theta}\left[\frac{\partial c_\theta}{\partial m} - 1\right]}_{=\frac{\partial l_\theta}{\partial m}} = 0.\end{aligned} \tag{C6}$$

I then assume separability of effort and consumption so that second cross-derivatives are zero. Separability has been widely used in the optimal unemployment literature to obtain tractable results (see e.g. Chetty (2008), Landais and Spinnewijn (2021)).

Assumption C3. $\frac{\partial^2 u_\theta}{\partial l \partial c} = 0$

This allows me to express C6 as

$$\frac{\partial^2 u_\theta}{\partial^2 c}\frac{\partial c_\theta}{\partial m} + \frac{1}{1-\tau_n}\frac{1}{\theta}\frac{\partial^2 u_\theta}{\partial^2 l}\frac{1}{1-\tau_n}\frac{1}{\theta}\left[\frac{\partial c_\theta}{\partial m} - 1\right] = 0.$$

Using the FOC in equation C3, we have that

$$\sigma_\theta^c \frac{\partial c_\theta}{\partial m} + \sigma_\theta^l \frac{1}{1-\tau_n} \frac{1}{\theta} \left[\frac{\partial c_\theta}{\partial m} - 1 \right] = 0,$$

where σ_θ^c and σ_θ^l are the coefficients of absolute risk aversion for consumption and labor effort. Re-arranging, this becomes

$$\frac{1}{1-\tau_n} \frac{1}{\theta} \frac{\sigma_\theta^l}{\sigma_\theta^c} = \frac{MPC_{z_\theta}}{1 - MPC_{z_\theta}}, \tag{C7}$$

where $MPC_{z_\theta} = \partial c_\theta / \partial m$ is the marginal propensity to consume out of unearned income at earnings z_θ. Combining this with the lower bound for the MRS between states θ and θ' in equation C5, we have that

$$MRS_{\theta,\theta'} \geq \frac{\sigma_\theta^c / \sigma_\theta^l}{\sigma_{\theta'}^c / \sigma_{\theta'}^l} \frac{O_{z_\theta}^{MPC}}{O_{z_{\theta'}}^{MPC}}, \tag{C8}$$

where $O_z^{MPC} = \frac{MPC_z}{1-MPC_z}$ is the odds of MPC at earnings level z.

Next, I make the assumption

Assumption C4. $\frac{\sigma_\theta^l/\sigma_\theta^c}{\sigma_{\theta'}^l/\sigma_{\theta'}^c} \leq 1$ *for all* θ *and* θ *such that* $\theta \leq \theta'$.

This together with equation C8 straight away implies that

$$MRS_{\theta,\theta'} \geq \frac{O_{z_\theta}^{MPC}}{O_{z_{\theta'}}^{MPC}}.$$

However, we are interested in estimating a lower bound for

$$\begin{aligned}
\beta(z) &= H(z) \frac{E(u_\theta^c | z_\theta < z) - E(u_\theta^c | z_\theta \geq z)}{E(u_\theta^c)} \\
&= H(z) \frac{E(\partial u_\theta / \partial c | z_\theta < z) - E(\partial u_\theta / \partial c | z_\theta \geq z)}{E(\partial u_\theta / \partial c)} \\
&= H(z) \frac{\delta(z) H(z) - 1}{\delta(z) H(z) + (1 - H(z))},
\end{aligned}$$

where $\delta(z) = E(\partial u_\theta / \partial c | z_\theta < z) / E(\partial u_\theta / \partial c | z_\theta \geq z)$. To get at this, note that as the optimality condition in C3 holds in each state θ, we can express the expected marginal utility of income above z as

$$E(u^c_\theta | z_\theta \geq z) = \sum_t \frac{\partial u_{\theta_t}}{\partial c} \frac{\mathbb{1}(z_{\theta_t} \geq z)}{1 - H_\theta(z)} = -\sum_t \frac{1}{1 - \tau_n(z_{\theta_t})} \frac{1}{\theta_t} \frac{\partial u_{\theta_t}}{\partial l} \frac{\mathbb{1}(z_{\theta_t} \geq z)}{1 - H_\theta(z)}$$
$$= -E\left(\frac{1}{1-\tau_n} \frac{1}{\theta} \middle| z_\theta \geq z \right) \kappa(z_\theta \geq z),$$

where

$$\kappa(z_\theta \geq z) = \sum_t \frac{\frac{\mathbb{1}(z_{\theta_t} \geq z)}{1-H_\theta(z)} \frac{1}{1-\tau_n(z_{\theta_t})} \frac{1}{\theta_t}}{E\left(\frac{1}{1-\tau_n} \frac{1}{\theta} \middle| z_\theta \geq z \right)} \frac{\partial u_{\theta_t}}{\partial l}$$

is the weighted average of $\frac{\partial u_{\theta_t}}{\partial l}$ above z, were the weights are given by $\frac{\mathbb{1}(z_{\theta_t} \geq z)}{1-H_\theta(z)} \frac{1}{1-\tau_n(z_{\theta_t})} \frac{1}{\theta_t}$.

Analogously, the expected marginal utility below z is

$$E(u^c_\theta | z_\theta < z) = -E\left(\frac{1}{1-\tau_n} \frac{1}{\theta} \middle| z_\theta < z \right) \kappa(z_\theta < z),$$

where

$$\kappa(z_\theta < z) = \sum_t \frac{\frac{\mathbb{1}(z_{\theta_t} < z)}{H_\theta(z)} \frac{1}{1-\tau_n(z_{\theta_t})} \frac{1}{\theta_t}}{E\left(\frac{1}{1-\tau(z)} \frac{1}{\theta} \middle| z_\theta < z \right)} \frac{\partial u_{\theta_t}}{\partial l}$$

is the weighted average of $\frac{\partial u_{\theta_t}}{\partial l}$ below z, were the weights are given by $\frac{\mathbb{1}(z_{\theta_t} < z)}{H_\theta(z)} \frac{1}{1-\tau_n(z_{\theta_t})} \frac{1}{\theta_t}$.

Then, $\delta(z) = E(u^c_\theta | z_\theta < z)/E(u^c_\theta | z_\theta \geq z)$ can be expressed as

$$\delta(z) = \frac{E\left(\frac{1}{1-\tau_n} \frac{1}{\theta} \middle| z_\theta < z \right)}{E\left(\frac{1}{1-\tau_n} \frac{1}{\theta} \middle| z_\theta \geq z \right)} \frac{\kappa(z_\theta < z)}{\kappa(z_\theta \geq z)}. \tag{C9}$$

As every θ for which $z_\theta < z$ is below any θ' for which $z_{\theta'} \geq z$, assumption C2 implies that $\frac{\kappa(z_\theta < z)}{\kappa(z_\theta \geq z)} \geq 1$. To see this, note that $\kappa(z_\theta < z)$ is a weighted average of $\partial u_\theta / \partial l$ for incomes below z and $\kappa(z_\theta \geq z)$ is a weighted average of $\partial u_\theta / \partial l$ for incomes above

z. As z_θ is increasing in θ, assumption C2 implies that $\partial u_\theta/\partial l \geq \partial u_{\theta'}/\partial l$ for each $z_\theta \leq z_{\theta'}$. Hence, all the points in the weighted average $\kappa(z_\theta < z)$ are above all the points in $\kappa(z_\theta \geq z)$, so that $\kappa(z_\theta < z) \geq \kappa(z_\theta \geq z)$ which implies the result.

Hence, assumption C2 and equation C9 imply that

$$\delta(z) \geq \frac{E\left(\frac{1}{1-\tau_n}\frac{1}{\theta}\middle| z_\theta < z\right)}{E\left(\frac{1}{1-\tau_n}\frac{1}{\theta}\middle| z_\theta \geq z\right)}. \tag{C10}$$

Next, taking the expected value of equation C7 below z gives

$$\sum_t \frac{1}{1-\tau_n(z_{\theta_t})}\frac{1}{\theta_t}\frac{\sigma^l_{\theta_t}}{\sigma^c_{\theta_t}}\frac{\mathbb{1}(z_{\theta_t} < z)}{H_\theta(z)} = \sum_t \frac{MPC_{z_\theta}}{1-MPC_{z_\theta}}\frac{\mathbb{1}(z_{\theta_t} < z)}{H_\theta(z)} \tag{C11}$$

This can be further expressed as

$$E\left(\frac{1}{1-\tau_n}\frac{1}{\theta}\middle| z_\theta < z\right)\psi(z_\theta < z) = E\left(\frac{MPC}{1-MPC}\middle| z_\theta < z\right), \tag{C12}$$

where

$$\psi(z_\theta < z) = \sum_t \frac{\frac{1}{1-\tau_n(z_{\theta_t})}\frac{1}{\theta_t}\frac{\mathbb{1}(z_{\theta_t}<z)}{H_\theta(z)}}{E\left(\frac{1}{1-\tau_n)}\frac{1}{\theta}\middle| z_\theta < z\right)}\frac{\sigma^l_{\theta_t}}{\sigma^c_{\theta_t}}$$

is a weighted average of $\frac{\sigma^l_\theta}{\sigma^c_\theta}$ in states with pre-tax earnings below z with the same weights as in $\kappa(z_\theta < z)$.

Analogously, for states with pre-tax earnings above z, we have

$$E\left(\frac{1}{1-\tau_n}\frac{1}{\theta}\middle| z_\theta \geq z\right)\psi(z_\theta \geq z) = E\left(\frac{MPC}{1-MPC}\middle| z_\theta \geq z\right), \tag{C13}$$

where

$$\psi(z_\theta \geq z) = \sum_t \frac{\frac{1}{1-\tau_n(z_{\theta_t})}\frac{1}{\theta_t}\frac{\mathbb{1}(z_{\theta_t}\geq z)}{H_\theta(z)}}{E\left(\frac{1}{1-\tau_n}\frac{1}{\theta}\middle| z_\theta \geq z\right)}\frac{\sigma^l_{\theta_t}}{\sigma^c_{\theta_t}}$$

is a weighted average of $\frac{\sigma^l_\theta}{\sigma^c_\theta}$ in states with pre-tax earnings above z with the same weights as in $\kappa(z_\theta \geq z)$.

Hence, we have that

$$\frac{E\left(\frac{1}{1-\tau_n}\frac{1}{\theta}\middle| z_\theta < z\right)}{E\left(\frac{1}{1-\tau_n}\frac{1}{\theta}\middle| z_\theta \geq z\right)}\frac{\psi(z_\theta < z)}{\psi(z_\theta \geq z)} = \frac{E\left(\frac{MPC}{1-MPC}\middle| z_\theta < z\right)}{E\left(\frac{MPC}{1-MPC}\middle| z_\theta \geq z\right)}$$

Assumption C4 then implies that $\frac{\psi(z_\theta<z)}{\psi(z_\theta\geq z)} \leq 1$ using a similar argument as for the result that $\frac{\kappa(z_\theta<z)}{\kappa(z_\theta\geq z)} \geq 1$. Hence, we get that

$$\frac{E\left(\frac{1}{1-\tau_n}\frac{1}{\theta}\middle| z_\theta < z\right)}{E\left(\frac{1}{1-\tau_n}\frac{1}{\theta}\middle| z_\theta \geq z\right)} \geq \frac{E\left(\frac{MPC}{1-MPC}\middle| z_\theta < z\right)}{E\left(\frac{MPC}{1-MPC}\middle| z_\theta \geq z\right)} \tag{C14}$$

Combining inequalities C10 and C17, we have

$$\delta(z) \geq \frac{E\left(\frac{MPC}{1-MPC}\middle| z_\theta < z\right)}{E\left(\frac{MPC}{1-MPC}\middle| z_\theta \geq z\right)}. \tag{C15}$$

Hence, we have proposition 1:

$$\beta(z) \geq H(z)\frac{E\left(\frac{MPC}{1-MPC}\middle| z_\theta < z\right) - E\left(\frac{MPC}{1-MPC}\middle| z_\theta \geq z\right)}{E\left(\frac{MPC}{1-MPC}\right)}. \tag{C16}$$

Proposition 2 follows from proposition 1 simply by using the budget constraint:

$$c_\theta = z_\theta - \tau_n[z_\theta - z_n] - T(z_n) + m_\theta.$$

From the budget constraint, we have that

$$MPC_{z_\theta} = \frac{\partial c_\theta}{\partial m} = (1-\tau_n)\frac{\partial z_\theta}{\partial m} + 1 = (1-\tau_n)MPE_{z_\theta} + 1,$$

where $MPE_{z_\theta} = \frac{\partial z_\theta}{\partial m}$ is the marginal propensity to earn out of unearned income at income level z_θ. Plugging this into equation C17 and simplifying, gives

$$\beta(z) \geq H(z)\frac{E\left(\frac{1/(1-\tau_n)+MPE}{MPE}\middle| z_\theta < z\right) - E\left(\frac{1/(1-\tau_n)+MPE}{MPE}\middle| z_\theta \geq z\right)}{E\left(\frac{1/(1-\tau_n)+MPE}{MPE}\right)}. \tag{C17}$$

D Appendix: Derivation of proposition 3

In this Appendix, I derive proposition 3:

Proposition 3. *Under assumption 1*

$$\beta(z) \approx H(z)\frac{E\left(1/c_\theta^{\gamma_\theta}|z_\theta < z\right) - E(1/c_\theta^{\gamma_\theta}|z_\theta \geq z)}{E(1/c_\theta^{\gamma_\theta})}.$$

First, consider the first order Taylor-approximation of the inverse of the marginal utility of income $\partial u(c,l,\theta)/\partial c$ at x:

$$\frac{1}{\partial u(c,l,\theta)/\partial c} \approx \frac{1}{\partial u(x,l,\theta)/\partial c} + \frac{\partial^2 u(x,l,\theta)/\partial^2 c - \partial^2 u(x,l,\theta)/\partial c\partial l \times \partial l/\partial c}{[\partial u(x,l,\theta)/\partial x]^2}[c-x]$$

Under assumption 1 and denoting $\partial u(c,l,\theta)/\partial c = u'(c)$ and $\partial^2 u(c,l,\theta)/\partial^2 c = u''(c)$, this can be expressed as

$$\frac{1}{u'(c)} \approx \frac{1}{u'(x)} - \frac{u''(x)}{u'(x)^2}[c-x].$$

Re-arranging this becomes

$$\frac{u'(c)-u'(x)}{u'(c)} \approx -\frac{u''(x)x}{u'(x)^2}\frac{x-c}{x} = -\gamma(x)\frac{c-x}{x}.$$

For c and x close to each other, $\frac{u'(c)-u'(x)}{u'(c)} \approx log(u'(c)) - log(u'(x))$ and $\frac{c-x}{x} \approx log(c) - log(x)$, so that we have

$$log(u'(c)) - log(u'(x)) \approx -\gamma(x)[log(c) - log(x)],$$

which implies that

$$u'(c) \approx \frac{x^{\gamma(x)}u'(x)}{c^{\gamma(x)}}.$$

Hence, as $\beta(z) = \left[E(u'(c)|z_\theta < z) - E(u'(c)|z_\theta \geq z)\right]/E(u'(c))$. We then have two possibilities: First, if $log(u'(x)) + \gamma(x)log(x) = 0$, we have that $u'(c) \approx 1/c^{\gamma(x)}$, which straight away implies proposition 3. Second, if $log(u'(x)) + \gamma(x)log(x) \neq 0$ we have that

$$\begin{aligned}
\beta(z) &= H(z)\frac{E(x^{\gamma(x)}u'(x)/c^{\gamma(x)}|z_\theta < z) - E(x^{\gamma(x)}u'(x)/c^{\gamma(x)}|z_\theta \geq z)}{E(x^{\gamma(x)}u'(x)/c^{\gamma(x)})} \\
&= H(z)\frac{x^{\gamma(x)}u'(x)E(1/c^{\gamma(x)}|z_\theta < z) - x^{\gamma(x)}u'(x)E(1/c^{\gamma(x)}|z_\theta \geq z)}{x^{\gamma(x)}u'(x)E(1/c^{\gamma(x)})} \\
&= H(z)\frac{E(1/c^{\gamma(x)}|z_\theta < z) - E(1/c^{\gamma(x)}|z_\theta \geq z)}{E(1/c^{\gamma(x)})}.
\end{aligned}$$

As we can choose the point of approximation x freely, one can choose $x = z$, this proves proposition 3.

E Appendix: Relationship to Baily-Chetty

The standard approach to characterizing the optimal level of social insurance in a two state model without redistributive preferences is the Baily-Chetty formula (Baily (1978); Chetty (2006)). In this appendix, I illustrate the connection between the

ABC-formula in that case (i.e. equation 12 without redistributive preferences, so that $g^I = g^R = 1$) and the Baily-Chetty formula. The difference between these two approaches arises from the way they treat the budget constraint: in the Baily-Chetty approach, the budget constraint is respected explicitly, whereas in the ABC-type formula it is not. Hence, I illustrate the connection between these two approaches by showing that the ABC-formula and the government budget constraint together imply the Baily-Chetty formula.

Consider the ABC-formula 12 without redistributive preferences, so that $g^R = g^I = 1$:

$$\frac{T_E - T_U}{c_E - c_U} = \frac{1}{\eta}\beta = \frac{1}{\eta}(1-h)\frac{u^c_U - u^c_E}{E(u^c)}, \tag{E1}$$

where $\beta = (1-h)\frac{u^c_U - u^c_E}{E(u^c)}$ by definition, and the budget constraint

$$hT_E + (1-h)T_U.$$

My goal is to show that these together imply the Baily-Chetty condition:

$$\varepsilon = h\frac{u^c_U - u^c_E}{u^c_E},$$

where $\varepsilon = \frac{d(1-h)}{d(-T_U)}\frac{-T_U}{1-h}$ is the elasticity of the unemployment rate $1-h$ with respect to the unemployment benefit $-T_U$.[17]

I will first derive an expression that relates $\eta = \frac{\partial(1-h)}{\partial(-T_U)}\frac{c_E - c_U}{h}$ and ε together. As noted above, the Baily-Chetty condition explicitly closes the government budget constraint. This implies that the change in T_E (dT_E) needs to be counteracted with a change in T_U (dT_U) such that the budget is in balance. Hence, h changes both

[17]Note that $\frac{d(1-h)}{d(-T_U)} \neq \frac{\partial(1-h)}{\partial(-T_U)}$ as also T_E changes in the Baily-Chetty setup so that there is no impact on the government budget.

because of a change in T_E and because of a change in T_U, so that

$$dh = \frac{\partial h}{\partial T_E} dT_E + \frac{\partial h}{\partial T_U} dT_U.$$

Due to the assumption of no income effects, we have that $\frac{\partial h}{\partial(c_E - c_U)} = -\frac{\partial h}{\partial T_E} = \frac{\partial h}{\partial T_U}$, so that

$$dh = \frac{\partial h}{\partial(c_E - c_U)}(dT_U - dT_E) = \eta \frac{h}{c_E - c_U}(dT_U - dT_E)$$

or

$$\eta = dh \frac{c_E - c_U}{h} \frac{1}{dT_U - dT_E}. \tag{E2}$$

Next, note that the budget constraint $hT_E + (1-h)T_U = 0$ implies that $dhT_E + hdT_E - dhT_U + (1-h)dT_U = 0$, or $dT_U - dT_E = \frac{dh}{h}(T_E - T_U) + \frac{dT_U}{h}$. Inserting this to equation E2, we have

$$\eta = dh(c_E - c_U) \frac{1}{dh(T_E - T_U) + dT_U}.$$

Re-arranging and noting that $dh = -d(1-h)$, we have that

$$\begin{aligned} \eta &= -\frac{d(1-h)}{d(-T_U)}(c_E - c_U) \frac{1}{\frac{d(1-h)}{d(-T_U)}(T_E - T_U) - 1} \\ &= -\varepsilon \frac{1-h}{(-T_U)}(c_E - c_U) \frac{1}{\varepsilon \frac{1-h}{(-T_U)}(T_E - T_U) - 1} \\ &= -\frac{c_E - c_U}{T_E - T_U} \frac{\varepsilon}{\varepsilon + \frac{1}{1-h} \frac{T_U}{T_E - T_U}}. \end{aligned}$$

Noting from the budget constraint that $T_E - T_U = \frac{-T_U}{h}$, this becomes

$$\eta = \frac{c_E - c_U}{T_E - T_U} \frac{\varepsilon}{\frac{h}{1-h} - \varepsilon}. \tag{E3}$$

Inserting this expression for η to the ABC-formula in equation E1 and re-arranging, we have that

$$
\begin{aligned}
\varepsilon &= \frac{h(u_U^c - u_E^c)}{E(u^c) - (1-h)(u_U^c - u_E^c))} \\
&= \frac{h(u_U^c - u_E^c)}{(1-h)u_U^c + hu_E^c - (1-h)u_U^c + (1-h)u_E^c} \\
&= h\frac{u_U^c - u_E^c}{u_E^c},
\end{aligned}
\tag{E4}
$$

where the second equality uses $E(u^c) = (1-h)u_U^c + hu_E^c$. Hence, we have the Baily-Chetty condition.