Designing a Basic Income: Lessons From the Optimal Income Tax Literature

Robin Boadway, Queen's University
Katherine Cuff, McMaster University

Date: January 2021

Author Note

The authors can be contacted at boadwayr@econ.queensu.ca and cuffk@mcmaster.ca.

We are grateful for comments received from the BC Basic Income Expert Committee members, David Green, Jonathan Rhys Kesselman, and Lindsay Tedds. Yiannis Kipouros and Barnabas Ney provided excellent research assistance.

Research paper commissioned by the Expert Panel on Basic Income, British Columbia. We gratefully acknowledge funding from the Government of British Columbia (spcs46008190052 and spsc46008190046) that helped support this research. All inferences, opinions, and conclusions drawn in this paper are those of the authors, and do not reflect the opinions or policies of the Data Innovation Program or the Province of British Columbia.

Abstract

The literature on optimal income taxation is vast, and it has grown rapidly in recent years. Our purpose here is to provide a broad overview of the literature and its implications for basic income. We show that the optimal income tax literature prescribes a basic income guarantee as an implicit component of the optimal tax structure. However, little attention is paid to practical tax design and institutional considerations. We briefly discuss such practical matters, including administrative body and structure, benefit reduction rates, federalism, financing, and work incentives.

Introduction

The optimal income tax literature studies optimal redistribution from a normative perspective. The objective of the government is to maximize a social welfare function that satisfies specified ethical properties. It can be contrasted with positive approaches to income redistribution that seek to explain how the political process determines government tax-transfer policy. Given the redistribution objectives built into the optimal income tax approach, it is not surprising that a basic income guarantee will characterize the solution. The form of the basic income depends on the particular assumptions made about the social welfare function and about the functioning of the economy.

The literature on optimal income taxation is vast, and it has grown rapidly in recent years. Our purpose here is to provide a broad overview of the literature and its implications for basic income. We first discuss the two standard models, which differ in their treatment of labour supply. In the *intensive-margin model* originating in Mirrlees (1971), individuals adjust the intensity of their labour supply in response to after-tax wages and incomes. In the *extensive-margin model* of Diamond (1980), individuals make discrete labour supply decisions such as labour market participation and job choice. These models share some common features, including a given distribution of workers by skill or productivity levels, a perfectly elastic demand for labour, and common social welfare functions. They also emphasize information as a basic constraint on government achieving its redistributive objectives.

We outline the tax-transfer structures obtained from these two approaches and stress their policy-relevant features, including those obtained from simulations. We summarize evidence in support of the social welfare functions adopted and about the behavioural responses to redistributive taxation. We discuss some consequences of relaxing important features of the models by allowing for heterogeneous preferences for leisure and for interdependent utility functions that take into consideration altruism, envy, and so on. We also discuss the consequences of recent findings in behavioural economics for tax-transfer policy. We discuss the effect of labour market considerations for optimal redistribution policy. This includes making wages endogenous and allowing for involuntary unemployment. These have implications for basic income, including the relative size of transfers given to the temporary and permanent unemployed, as well as those who choose not to work or are unable to work. The literature on this is still in its infancy, so our discussion is mainly suggestive. We finish with some practical lessons learned about the design and implementation of a basic income guarantee for Canada.

Architecture of Basic Income

Basic income generally refers to a system of income transfers that ensures that all individuals attain some level of income, and it can take different forms. We classify basic income systems according to three different types, which we can relate to the standard literature on optimal income taxation. The latter focuses on both the size of income transfers and their optimal financing by income taxation. In what follows, we consider only income taxes, although governments use other taxes for revenue-raising, such as sales and excise taxes, business taxes, and property taxes. We also focus on individuals as the taxpaying unit, setting aside issues that arise with family or household taxation.

Basic Income Classification

Universal Basic Income

A *universal basic income* is the least restrictive form of basic income and has been prominently proposed by Van Parijs (1995) and Standing (2017).[1] A universal basic income is provided in the same amount to all individuals regardless of their circumstances or behaviour. Financing can come from various sources, including general tax revenues or a dedicated source of revenues like a natural resource tax. Proponents of universal basic income often emphasize the benefit of a non-discriminatory income transfer and focus less on its finance. The financial cost can be substantial: a universal transfer to all individuals of, say, \$20,000 per year naturally costs \$20,000 per capita to finance, although the financing would necessarily not be equal per capita. We discuss below how a universal basic income can be interpreted as a component of the income tax structure, including the optimal income tax system.

Targeted Basic Income

A *targeted basic income* means that different target groups receive transfers that differ in size. Targeting may be by age (child, working-age adult, senior), family composition (single-person versus multi-person families), or health status (able versus disabled). Targeting—or tagging—relies on identifiable characteristics to determine the allocation of transfers. Targeted transfers require a

[1] The idea of a universal basic income has a long history. Well-known proponents include, alphabetically, Atkinson (2015), Cole (1935), de Caritat (1795/1988), Fourier (1836/1967), Friedman (1962), Galbraith (1986), George (1871), King (1967), Meade (1936), Mill (1849/1987), More (1516/1963), Paine (1796), Reich (2016), Russell (1918), Simon (2000), Tinbergen (1975), Tobin (1966), and von Hayek (1979).

choice of the relevant target groups as well as a judgment about the relative size of the transfer to different groups. Within a given target group, all persons receive the same transfer, regardless of their behaviour or their circumstances. As with a universal basic income, financial costs can be substantial.

Conditional Basic Income

A *conditional basic income* bases the size of the transfer on need, deservedness, or some action of the recipient. An individual's income is commonly taken to be a measure of need, and this has the merit of being consistent with the income tax system (Segal, 2016; Boadway et al., 2018a, 2018b). Other indicators of need could be contemplated, such as asset wealth or consumption, though these may be more difficult to administer and could lead to more intrusion into individual privacy given the informational requirements. Atkinson (2015) and Osberg (2018) have proposed that basic income transfers be contingent on some form of civic participation, in part to make it more politically and socially acceptable. Civic participation can include participating (or seeking to participate) in the labour market, but could also include operating a business, training, education, volunteer work, and caregiving. One would need some means of determining if the criterion of civic participation is satisfied, and this too could be intrusive.

Making basic income conditional on income or participation leads to concerns about incentive effects, particularly work incentives. Work incentive effects are central to the optimal income tax literature and will be relevant for our discussion below. They depend on both income and price effects. A universal basic income has no price effects, as it does not affect the relative return to working, but it still affects work incentives through income effects. A conditional basic income will have both price and income effects, and the price effect on the relative return to working will be greater the more basic income transfers decline with income.

Guaranteed Basic Income

A *guaranteed basic income* refers to the idea that no person—perhaps, of a given type—will be left without at least some minimum level of income. A guaranteed basic income could take any of the above forms: universal with the same guaranteed minimum income for all individuals, or a targeted or conditional guaranteed income. Its emphasis is on ensuring that those with the least income attain some guaranteed minimum. As we argue below, there is a fundamental similarity between universal guaranteed basic income and income-contingent guaranteed basic income once one takes the financing into account.

Basic Income in Optimal Income Taxation

The standard approach to optimal income taxation studies how the tax-transfer system should be designed to achieve some social objective given constraints the government faces, including restrictions on its policy instruments, imperfect information, and revenue requirements. Before discussing the optimal income taxation approach (see "Two Approaches to Optimal Income Taxation"), we briefly outline three versions of income taxes and transfers studied in the optimal income literature. All three versions yield a guaranteed basic income conditional on labour earnings. The first is the *linear progressive income tax*, or flat tax, where the government is restricted to combining a single income tax rate with a credit or exemption level. This can be generalized in a straightforward way to a *piecewise linear progressive income tax*, where, rather than imposing a uniform tax rate, the government can choose multiple tax brackets, each with its own associated tax rate. The third approach yields a *nonlinear income tax* defined over all income levels. Each version contains a guaranteed income level that applies to the lowest-income individuals.

Linear Progressive Income Tax

Suppose the government is restricted to choosing a two-parameter income tax function, $T(y) = ty - b$, where y is taxable income, t is a constant marginal tax rate, and b is a credit available to all taxpayers. (We ignore complications arising from family size, disability, and other factors that might lead to some targeting.) Individuals differ in their wage rates, and income is the product of the wage rate and amount of labour supplied, $y = w\ell$. Those with higher wage rates earn higher levels of income for a given amount of labour supplied. The optimal linear progressive income tax is the combination of t and b that maximizes social welfare, given individual behaviour and government budget balance. Assuming that social welfare maximization entails redistributing from higher- to lower-income persons, the credit b will be positive, and the marginal tax rate t will be between zero and unity. The exact sizes of t and b will depend on the distribution of wage rates; the behavioural responses of individuals to taxation, which depends on individual preferences; and the inequality aversion of the social welfare function.

Under a linear progressive income tax, all individuals face a budget constraint:

$$c = y - T(y) = (1-t)y + b \tag{1}$$

where c is disposable income, or consumption. Figure 1 depicts this budget constraint, which is labelled NIT to reflect the fact that a linear progressive income tax is equivalent to a negative income tax originally proposed by Friedman (1962). Tax liability at income level y, or $T(y)$, is the vertical distance between the no-tax budget line (the 45° line) and the after-tax budget line and is equal to zero at income y_0. Low-w individuals choose low income levels and have negative tax liabilities, while high-w individuals pay positive taxes. The tax system is progressive by design: the average tax rate, which is total tax liability divided by earned income, $T(y)/y$, is negative for low-income earners and increasing in income.

Three observations should be made about the linear progressive tax in Figure 1. First, for the budget constraint NIT to apply, negative tax liabilities must be refunded to low-income persons. In the absence of refundability, the tax schedule would follow the 45° line for those with incomes below the cut-off level y_0, and there would be no transfers in the income tax system. The average tax rate is then increasing only above the income cut-off and is zero below.

Second, the tax credit b is equivalent to a basic income guarantee. No one obtains an after-tax income of less than b. Those with zero income obtain a basic income transfer of b. The level of the basic income guarantee b that is compatible with government budget balance depends on the tax rate t.

Third, the assumption of a uniform tax rate restricts the size of the basic income guarantee b. To show this, the dashed line in Figure 1 depicts a budget constraint if the government imposes a higher tax rate for those below the cut-off level and a lower one for those above, such that overall budget balance is satisfied. The higher tax rate on incomes below the cut-off level induces these lower-income earners to reduce their labour supply and income. Suppose this reduction is not large enough to reduce the amount of the transfer paid to them. Less tax revenue then needs to be raised from income earners above the cut-off level, so the tax rate above the cut-off can be reduced. This allows both the basic income guarantee to increase to b^+ and the tax rate on higher-income individuals to fall while maintaining budget balance. Of course, if the labour supply of low-income persons is highly responsive, an increase in their marginal tax rate could increase the transfers they receive, so this analysis would not follow. The possibility that a higher basic income might be financed with lower tax revenues by increasing the marginal tax rate at the bottom is important in what follows.

Piecewise Linear Income Tax

A more general version of the dashed budget constraint in Figure 1 is the piecewise linear income tax, which allows for an arbitrary number of tax brackets and associated tax rates. Figure 2 depicts one simple extension of Figure 1 with four tax brackets. The first tax bracket applies for those below the break-even income level y_0 and has a relatively high marginal tax rate. The three tax brackets for persons above the break-even level are defined by the income ranges y_0 to y_1, y_1 to y_2, and greater than y_2. In this example, tax rates are increasing as one moves to higher tax brackets above the break-even level. The guaranteed basic income is again labelled b and applies to those with zero income. All individuals who choose to work have at least that level of after-tax income. As in the linear progressive income tax case, this assumes that negative tax liabilities are fully refunded. The marginal tax rate on incomes less than or equal to y_0 in the first tax bracket is relatively high. This facilitates a high basic income guarantee b, assuming low-income persons do not reduce their earnings enough to warrant a higher transfer. This marginal tax rate below the break-even income level is often known as the *clawback rate* or the *benefit-reduction rate*. After-tax income is reduced at this rate for each dollar of income earned.

While a higher marginal income tax rate applying at low incomes might allow a higher basic income guarantee, it also increases the disincentive to work. One way to encourage labour force participation is to reduce the marginal tax rate at very low income levels. The dashed line in Figure 2, which retains the piecewise linear form, depicts this. The marginal tax rate is zero for an initial amount of income and then becomes positive. Individuals retain the full basic income b for some initial amount of income that they earn, and then the income transfer begins to be clawed back. Financing this zero tax bracket requires raising tax rates at higher income levels. The marginal tax rate for those at low income levels could even be negative, as in the Canada Workers Benefit (CWB) and the Earned Income Tax Credit (EITC) in the United States.

The overall tax structure depicted in Figure 2 has a similar form to that found in Canada and other countries. There is a critical difference, however. Apart from refundable tax credits, transfers made to persons at low income levels are delivered largely by social assistance rather than through the tax system. These transfers typically have relatively high effective marginal tax rates and are administered by social workers as gatekeepers rather than by the tax administration based on self-reporting of taxpayers. Commonly advocated conditional basic income systems would deliver transfers through the income tax system using a self-reporting mechanism (Stevens & Simpson, 2017; Boadway et al., 2018b; Koebel & Pohler, 2019), which corresponds with the approach of the optimal income tax literature. Overall, the limited literature on the optimal piecewise linear tax system offers few general properties of the ideal tax structure (Strawczynski, 1988; Sheshinski, 1989; Slemrod et al., 1994; Apps et al., 2011; Bastani et al., 2019). Much of the optimal taxation literature focuses on either uniform linear or nonlinear income tax systems.

Optimal Nonlinear Income Tax

In Mirrlees' (1971) seminal paper on optimal income taxation (summarized in Atkinson and Stiglitz, 1980; Tuomala, 1990, 2016; and Myles, 1995), the government implements a nonlinear income tax function $T(y)$, where $T'(y)$ is the marginal income tax rate at income level y. Figure 3 depicts a typical budget constraint under such a tax where the slope of the budget constraint at income level y is $1-T'(y)$. In this figure, the structure of marginal tax rates is inverse U-shaped. The marginal tax rate is positive, relatively high and declining at low income levels, and then increasing at higher income levels. Once again, the tax schedule yields negative total tax liabilities at low income levels, and a basic income guarantee of b applies to individuals who choose to be voluntarily unemployed.

In the Mirrlees model, labour supply varies along the intensive margin, so individuals can vary their labour supply continuously from $\ell = 0$. We refer to this as the *intensive-margin labour supply approach*. Diamond (1980) and Saez (2002) assume, on the contrary, that labour supply involves a discrete decision. In the extreme, individuals decide only whether to participate in the labour force, and if they do, they obtain a job suited to their skills, work a fixed amount of time, and earn a job-specific income. In this case, referred to as the *extensive-margin labour supply approach*, Saez shows that the optimal nonlinear income tax policy typically involves a negative participation tax at the bottom of the income distribution, as shown in Figure 4. Marginal tax rates are initially very low, possibly negative, to encourage participation, and then follow a more standard pattern. For those not working, the basic income guarantee b applies, and everyone working obtains a net income greater than b.

Equivalence of Universal Basic Income and Income-Contingent Basic Income

It is a common misperception that since a universal basic income is more costly than an income-conditioned basic income, it is financially very challenging. However, this fails to take account of the methods of financing each. It is straightforward to show that once tax financing is taken into account, a universal basic income can be replicated by an income-contingent basic income, with one proviso related to the timing of the payments.[2] To see this, consider the linear progressive income tax of Figure 1. The tax liability of an individual with income y is $T(y) = ty - b$. Suppose a universal basic income of b is paid to all individuals and is financed by a proportional income tax. The net benefit to a person with income of y will be $b - ty$, which is precisely the net tax liability under a negative income tax.

Alternatively, the same outcome can be achieved by implementing a conditional basic income using a refundable tax credit mechanism combined with a proportional income tax. The refundable tax credit consists of a basic amount b subject to a taxback rate of t, which disappears at y_0 where $b - ty_0 = 0$. Beyond income level y_0, the standard negative income tax applies. This refundable tax credit mechanism can also implement the dashed budget constraint of Figure 1. In this case, the taxback rate applying to the refundable tax credit is larger than that applying to taxpayers with positive income taxes owing. These equivalences between a universal basic income and a negative income tax with refundability also apply with the piecewise linear income tax of Figure 2 and the nonlinear income tax of Figure 3.

Although the net effects of income tax–financed universal and income-conditioned basic income are the same, the timing of their payments differ. Universal basic income b is paid up front and is financed by subsequent income tax levies, ty. In the case of an income-contingent basic income, only net benefits $b - ty$ are paid to each individual after tax liabilities have been calculated. Thus, under a universal basic income, the government must finance the transfer up front and be reimbursed when taxes are paid. Since the same outcome is ultimately achieved, this upfront financing can be avoided by adopting the income-contingent version. This is also consistent with what would be prescribed by the optimal income tax literature, and in what follows, we assume that basic income takes the income-contingent type.

Two Approaches to Optimal Income Taxation

In this section, we review the two classic approaches to optimal income taxation associated with Mirrlees (1971) and Diamond (1980), respectively. In both approaches, individuals differ in an innate characteristic called productivity that determines the income they are able to earn. In the Mirrlees approach, individuals have the same preferences for goods and labour or leisure and can

[2] This equivalence has been noted previously by Kesselman and Garfinkel (1978), who compare the efficiency properties of a basic income credit tax system (where a single payment is made to all individuals and taxed back at the same rate) and a negative income tax system.

vary the intensity with which they supply labour. The government cannot observe productivity or wage rates, which individuals cannot influence, but can observe incomes, which they can choose. With perfect information, the government could condition individual lump-sum transfers and taxes on productivity and achieve any feasible distribution of income without any efficiency losses. In contrast, if the government can observe only incomes, individuals can choose their incomes to partially frustrate the government's desire for income redistribution. Formally, the government faces an incentive constraint that precludes it from redistributing income so much that higher-productivity individuals prefer to choose incomes corresponding with those of lesser productivity.

Diamond instead assumes that hours of work are fixed. Individuals differ in the income they will earn if they work and decide only whether to participate in the labour force. Individuals differ in their preferences for non-participation, and the government cannot observe those preferences. If the government imposes a higher tax on individuals of a given income, some of them will choose not to participate, and that constitutes the efficiency consequences of income taxation. In both the Mirrlees intensive-margin approach and the Diamond extensive-margin approach, individuals face perfectly elastic demands for their labour, so wage rates are fixed, and all increases in labour supply find employment.

The government chooses a tax structure to maximize a social welfare function that captures society's equity values, and it may be constrained in its choice of a tax function. In both approaches, optimal taxes are a trade-off of efficiency versus equity concerns and take into consideration the distribution of skills, individual preferences, and the form of the social welfare function. We consider the intensive- and extensive-margin approaches in turn. In each case, we emphasize the treatment of individuals at the bottom of the skill distribution, since these are the ones for whom the basic income guarantee is relevant.

Intensive-Margin Labour Supply Approach

The intensive-margin approach adopts the model used by Mirrlees (1971). The economy is populated by individuals who differ in their given productivities, which are reflected in their wage rate w. Wages are distributed by a distribution function $F(w)$ on $[\underline{w}, \overline{w}]$ with $\underline{w} \geq 0$ and density $f(w)$. The population is normalized to unity, so $F(\overline{w}) = 1$. Individuals of type w choose labour supply $\ell(w)$ and earn income of $y(w) = w\ell(w)$, also referred to as their *effective labour supply*. They pay taxes according to the income tax function $T(y)$, leaving $c(w) = y(w) - T\big(y(w)\big)$ as consumption. All individuals have the same utility function $u(c, \ell)$, which is increasing in c, decreasing in ℓ, and strictly concave.

The production side of the economy is particularly simple. Output is linear in effective labour supplied, so total output is $Y = \int_w w\ell(w)f(w)dw$. This implies that the demand for labour is perfectly elastic: if a type-w individual increases labour supply, the additional labour supplied will be employed at the wage w. This in turn implies that all responses to taxation come from the labour supply side of the economy. We return to alternative labour demand environments later.

The government chooses its tax function to maximize an additive social welfare function based on a vector of individual utilities $\boldsymbol{u}(\cdot)$:

$$\mathscr{W}\left(\boldsymbol{u}(\cdot)\right) = \int_{\underline{w}}^{\overline{w}} W\left(u(c(w), \ell(w))\right) f(w)dw = \int_{\underline{w}}^{\overline{w}} W\left(u\left(c(w), \frac{y(w)}{w}\right)\right) f(w)dw \tag{2}$$

where $W\left(u(c(w), \ell(w))\right)$ is a social utility function that is increasing in utility and concave, and $\mathscr{W}\left(\boldsymbol{u}\right)$ is symmetric in individual utilities. The assumption of concavity implies that the government has non-negative aversion to utility inequality. If social utility $W\left(u(\cdot)\right)$ is strictly concave, marginal social utility $W'\left(u(\cdot)\right)$ will be decreasing in individual utility. A redistribution of income from high- to low-utility persons increases social welfare. The degree of concavity of the social welfare function represents the weight put on equity by the government—that is, the *aversion to utility inequality*. In the *utilitarian* case, the social utility function is linear, and social welfare is linear in individuals' utilities. Since individual utilities are strictly concave in income, the government will still want to redistribute income from high- to low-income persons. The extreme case of inequality aversion is the *maximin* case, where the government only cares about the lowest-utility person. The government also faces a budget constraint of the form:

$$\int_{\underline{w}}^{\overline{w}} T\left(y(w)\right) f(w)dw = R \tag{3}$$

where $R \geqslant 0$ is the revenue required to finance given government expenditures.

We distinguish between two alternative forms of the government income tax function $T(y)$. First, we assume that $T(y)$ is linear progressive so consists of a constant marginal tax rate and a fixed tax credit for which all taxpayers are eligible, as in Figure 1. Second, we assume the optimal income tax function is nonlinear. In both cases, we assume that the government can observe individuals' incomes, but not their labour supplies or their wage rates. This assumption is motivated by the fact that in actual tax systems, individuals report their incomes to the tax authority. We assume that they do so truthfully. This is a strong assumption and requires that the tax administration can enforce truthful reporting by its system of audits and penalties.

Optimal Linear Income Taxation

Following Sheshinski (1972), suppose the government is restricted to a linear progressive income tax of the form $T(y) = ty - b$. To facilitate the government's problem, individual utility is written in the indirect form as a function of the tax parameters (t,b), where the indirect utility function is defined by

$$v(t,b;w) \equiv \max_{\{c,\ell\}} u(c,\ell) \quad \text{s.t.} \quad c = (1-t)w\ell + b. \tag{4}$$

The indirect utility function is useful for analyzing tax policy because it takes into account individual behavioural responses to tax changes. The value of the Lagrange multiplier on the budget constraint in problem (4) is called the *marginal utility of income*, and given the linear progressive tax system is equivalent to $\partial v(t,b;w)/\partial b$. The individual's marginal utility of income reflects the increase in the maximized level of utility with an additional dollar in after-tax income or, equivalently, a marginal increase in b. Indirect utility is increasing in wages, and the marginal utility of income is non-increasing in wages. Higher-wage individuals value an additional dollar in after-tax income less than lower-wage individuals.

The government chooses its tax parameters to maximize social welfare subject to its budget constraint:

$$\max_{\{t,b\}} \int_{\underline{w}}^{\overline{w}} W\big(v(t,b;w)\big) f(w)dw \quad \text{s.t.} \quad \int_{\underline{w}}^{\overline{w}} tw\ell(t,b;w)f(w)dw = b + R \tag{5}$$

where we have used $\int_{\underline{w}}^{\overline{w}} bf(w)dw = b$ since population size is normalized to unity. The solution to this problem will yield optimal values of b and t, and therefore the optimal negative income tax scheme corresponding to Figure 1.

To facilitate the interpretation of the solution to this problem, it is useful to introduce two definitions. One is the *marginal social utility of income* for a type-w individual, which is the change in social welfare from an increase in the after-tax income or consumption of a type-w person:

$$\beta(w) \equiv \frac{\partial W\big(v(t,b;w)\big)}{\partial v(t,b;w)} \cdot \frac{\partial v(t,b;w)}{\partial b}. \tag{6}$$

The marginal social utility of income is the product of two terms. The first is the change in social welfare from an increase in type-w's utility. The second is type-w's marginal utility of income, as defined above. Since the social utility function is concave and individual utilities are strictly concave, $\beta(w)$ is decreasing in w: an increment of after-tax income or consumption is more socially valuable for lower-wage persons. The marginal social utility of income is defined for a given tax rate.

Next, the *net social marginal value of income* extends $\beta(w)$ by taking account of the fact that giving an increment of after-tax income to an individual has indirect effects on government revenue by changing the amount of income the individual will earn. It is evaluated in terms of government revenue and is defined as:

$$\alpha(w) \equiv \frac{\beta(w)}{\lambda} + t\frac{\partial y(t,b;w)}{\partial b} \tag{7}$$

where λ is the shadow value of government revenue, or the value of the Lagrange multiplier on the government revenue constraint in the government's optimal tax problem, (5). The second term in (7) represents an indirect revenue effect arising from the behavioural response of the individual to an increase in non-labour income. Assuming leisure is normal, this second term is negative and larger for higher-w persons. Then, $\alpha(w)$ will also be decreasing in w.

Armed with these definitions, the first-order conditions to the government's problem for t and b reduce to the following:

$$\frac{t}{1-t} = -\frac{\int_{\underline{w}}^{\overline{w}} \big(\alpha(w)-1\big)w\ell f(w)dw}{\int_{\underline{w}}^{\overline{w}} w\ell\varepsilon_{\ell} f(w)dw} = -\frac{\text{Cov}[\alpha,y]}{\int_{\underline{w}}^{\overline{w}} y\varepsilon_{\ell} f(w)dw} = \frac{\text{equity}}{\text{efficiency}} \tag{8}$$

$$\mathbb{E}[\alpha] = 1 \tag{9}$$

where ε_{ℓ} is the compensated elasticity of labour supply with respect to the after-tax wage rate, and $y = w\ell$. The optimal tax rate in (8) is the ratio of an equity and an efficiency effect. The tax rate is inversely related to an aggregate of compensated elasticities of labour supply weighted by total incomes at each wage level. This is the distortionary or efficiency effect of the income tax. The tax rate is increasing in the covariance of the net social marginal value of income with income. The more rapidly $\alpha(w)$ declines with w (i.e., the more concave are the social utility and individual utility functions), and the more disbursed is income, the higher the tax rate. Eq. (8) captures in a standard way the equity-efficiency trade-off in choosing the tax rate and will be a useful basis for comparison with later results.

The optimal lump-sum transfer b, or basic income guarantee, satisfies (9). The left-hand side represents the average net social marginal benefit of an incremental increase in the basic income, and the right-hand side represents the marginal cost of this increase (since the population is normalized to one). The optimal basic income guarantee ensures these two expressions are equal.

As (8) and (9) indicate, the optimal values of t and b depend on the underlying wage distribution, the compensated elasticity of labour supply, the revenue requirement, and the social welfare function.

Under the special case of no income effects and identical constant elasticity of labour supply, Tanninen et al. (2019) show analytically that when the wage distribution is unbounded and follows a Pareto distribution, and the social welfare function is of the constant aversion-to-utility-inequality form, the optimal t and b are increasing in the degree of aversion to utility inequality. They also show that the optimal basic income b is decreasing in the government revenue requirement. The more revenue is required, the larger is λ and the higher is the tax rate t, but less of the tax revenue is available for the lump-sum transfer b.

Two other special cases are relevant for the choice of a basic income. One is where the social welfare function takes the maximin form. In the second, there are multiple commodities, each with its own commodity tax rate. We consider them in turn.

With a maximin social welfare function, the government maximizes the welfare of the lowest-utility person, who is the one with the lowest wage rate $\underline{w}$, so $\beta(w) = 0$ for all $w > \underline{w}$. This case also yields the largest feasible basic income guarantee b when the government is restricted to a linear progressive income tax. Suppose for simplicity that $\underline{w} = 0$, so the lowest-wage person earns no income. Then, the welfare of the least well-off person is maximized when the government chooses t to maximize b from the government budget constraint in (5). This yields the *tax-revenue maximizing tax rate* given by (8), with $\alpha(w) = t\partial y(t,b;w)/\partial b$ for all $w > \underline{w}$. Without income effects and if all individuals have the same constant compensated elasticity of labour supply ε, the maximin optimal tax rate is $t/(1-t) = 1/\varepsilon$, which is less than one. The government is constrained in its ability to raise tax revenue only by the behavioural responses of individuals. Provided the lowest-wage individuals are not working, the government does not need to be concerned about how the tax rate affects these individuals, since they are not paying any income taxes.

If $\underline{w} > 0$ so $y(\underline{w}) > 0$, then a government with a maximin objective must take into account how the tax rate impacts the lowest-wage individuals. The optimal tax rate with no income effects and constant elasticity in this case will be

$$\frac{t}{1-t} = \frac{1 - \frac{y(\underline{w})}{\mathbb{E}(y)}}{\varepsilon}$$

where the term in the numerator reflects the equity concerns for the lowest-wage individuals and will necessarily be less than one, so the optimal tax rate will be lower than the tax revenue–maximizing rate. The higher the income of the lowest-wage individuals relative to the average income, the lower the optimal tax rate.

Turn now to the second case, where we allow for multiple commodities, each with its own commodity tax rate. Suppose there are n commodities, $\boldsymbol{x} = (x_1, \cdots, x_i, \cdots, x_n)$, and write utility as $u(\boldsymbol{x}, \ell)$. The government can now impose a linear progressive income tax with parameters t and b and proportional commodity taxes τ_i. If commodity quantities are measured such that producer prices equal unity, the individual's budget constraint can be written as:

$$\sum_{i=1}^{n}(1+\tau_i)x_i = (1-t)y+b. \tag{10}$$

Deaton (1979) showed that if the government is restricted to a linear progressive income tax, commodity taxes should not be differentiated if individual preferences are (a) weakly separable in commodities and leisure and (b) quasi-homothetic in commodities (i.e., homothetic to any point). That is, Engel curves for all goods are linear and of the same slope for all individuals. If the government imposes commodity taxes, they should be uniform, so they add little to what is achieved when the government relies solely on an income tax. Hellwig (2009, 2010) generalized this result. He showed that if the Deaton conditions apply, starting from a differentiated commodity tax system and an arbitrary linear progressive income tax, a Pareto-improving tax reform can be designed that makes commodity taxes uniform and adjusts the parameters t and b of the income tax system appropriately.

Note that even though uniform commodity taxation and proportional income taxation are equivalent in this case, the decision to impose commodity taxes has implications for the basic income guarantee b. To see this, consider the budget constraint (10) when commodity taxes are uniform, so $\tau_i = \tau$ for all i:

$$\sum_{i=1}^{n}(1+\tau)x_i = (1-t)y+b.$$

If the government eliminates commodity taxes and adjusts the income tax system so the budget constraint still applies, the latter may be written as:

$$\sum_{i=1}^{n} x_i = \frac{1-t}{1+\tau}y + \frac{b}{1+\tau} = (1-t')y+b'$$

where $t' = (t+\tau)/(1+\tau)$ and $b' = b/(1+\tau)$. That is, when commodity taxes are used, the basic income level b must be adjusted upward to maintain the same purchasing power. This will apply to more general tax systems as well.

When the Deaton conditions do not apply, it is desirable to differentiate commodity taxes.

Even in this case, the *level* of commodity taxes can be chosen arbitrarily. If this level is positive, the lump-sum transfer b has to adjust upward.

Optimal Nonlinear Income Taxation: The Standard Analysis

Suppose now the government is able to use a nonlinear income tax $T(y)$, so a type-w individual's budget constraint is $c(w) = y(w) - T\big(y(w)\big)$. In the standard optimal income tax analysis beginning with Mirrlees (1971), the government directly chooses $y(w)$ ($= w\ell(w)$) and $c(w)$, and the structure of $T\big(y(w)\big)$ is inferred from the first-order conditions to the government's problem. The problem is highly technical, so we focus mainly on stating the problem and discussing the results. In choosing $y(w)$ and $c(w)$, the government's choices must satisfy the individuals' utility maximization behaviour, and the latter is represented by a set of incentive-compatibility, or implementability, conditions reflecting the individuals' first-order conditions. Thus, the government chooses $y(w)$ and $c(w)$ to maximize (2), subject to incentive-compatibility conditions for each individual type and the budget constraint $\int_w \big(y(w) - c(w)\big) f(w)dw = R$. Suppose, to simplify matters, that individual utility functions are additively separable in consumption c and labour $\ell = y/w$: $u(c, y/w) = \mu(c) - h(y/w)$. The expression for the optimal marginal income tax rate for a type-w individual obtained from the first-order conditions of the government's problem is:

$$\frac{T'\big(y(w)\big)}{1 - T'\big(y(w)\big)} = \mu'\big(c(w)\big) \int_w^{\overline{w}} \frac{1}{\mu'\big(c(\tilde{w})\big)} \left(1 - \frac{W'(\tilde{w})\mu'\big(c(\tilde{w})\big)}{\lambda}\right) dF(\tilde{w}) \frac{1 + \ell h''/h'}{w f(w)} \tag{11}$$

where λ is the Lagrange multiplier on the government's budget constraint and $W'(w) \equiv \partial W(v(w))/\partial v(w)$, and $v(w)$ is the maximized level of individual utility. The term $W'(\tilde{w})\mu'\big(c(\tilde{w})\big)/\lambda$ is the marginal social welfare weight for a type-$\tilde{w}$ person, analogous to $\beta(w)$ in the linear case, and reflects the increase in social welfare from increasing $c(\tilde{w})$ by an incremental amount. It will also be decreasing in the wage rate, given our assumptions on social welfare and individual preferences.

The intuition for the marginal tax rate in (11) can be seen by assuming that the utility function is quasilinear in consumption, so $u(c, y/w) = c - h(y/w)$, implying that $\mu'\big(c(w)\big) = 1$. In this case, there are no income effects on labour supply, so the compensated and uncompensated elasticities of labour supply are the same and satisfy $\ell h''/h' = 1/\varepsilon_\ell$, so (11) can be written:

$$\frac{T'\big(y(w)\big)}{1 - T'\big(y(w)\big)} = \underbrace{\int_w^{\overline{w}} \left(1 - \frac{W'(\tilde{w})}{\lambda}\right) dF(\tilde{w})}_{\text{equity}} \cdot \underbrace{\frac{1 + \varepsilon_\ell}{\varepsilon_\ell w f(w)}}_{\text{efficiency}}. \tag{12}$$

As in the linear progressive tax case of (8), the efficiency term is inversely related to the elasticity of labour supply and to the size of the wage base $wf(w)$. It reflects the reduction in the tax base as a result of an increase in the marginal tax rate at wage rate w. The equity term reflects the effect on social welfare of an increase in the marginal tax rate. An increase in $T'(w)$ causes tax liabilities to rise by an equal lump-sum amount for all individuals with wages above w. This simultaneously reduces social welfare measured in terms of government revenue by W'/λ and raises government revenue by a unit. The equity term can rise or fall with w: the range of the integral in (12) falls with w, but the average size of the term in the integral rises because marginal social utility $W'(w)$ is decreasing in w. Therefore, while $T'\big(y(w)\big)$ is non-negative, it can rise or fall with w so can take on many different patterns.

Note that the equity term is zero at both $\underline{w}$ and $\overline{w}$, so optimal marginal tax rates at the top and bottom are zero. The significance of the zero marginal tax rate at the top can be discounted since it applies only at the very top, and not immediately below the top. Moreover, the average tax rate at the top can be large even if the marginal tax rate is low.

At the bottom, the zero marginal tax rate is of limited relevance. If there is bunching at the bottom, so many wage types select the same bundle of y and c (for example, $y = 0$), the marginal tax rate will be positive. To see this, suppose that $\underline{w} = 0$ so the lowest-wage individual supplies zero labour and receives $-T(0)$ in consumption. Further, suppose there is some $w_0 > 0$ such that $w_0\big(1 - T'(0)\big)\mu'\big(-T(0)\big) - h'(0) = 0$, which implies that all individuals with $w \leq w_0$ will also choose to supply zero labour and there will be $F(w_0)$ individuals earning zero income. Piketty and Saez (2013) show that the optimal marginal tax rate on these non-working individuals is given by

$$\frac{T'(0)}{1-T'(0)} = \frac{g_0 - 1}{\varepsilon_0}, \quad \text{where} \quad g_0 = \frac{W'(w_0)\mu'(-T(0))}{\lambda}, \; \varepsilon_0 = -\frac{1-T'(0)}{F(w_0)}\frac{dF(w_0)}{d(1-T'(0))} \qquad (13)$$

so g_0 is the marginal social welfare weight on the non-working individuals in terms of government revenue, and ε_0 is the elasticity of the number of non-working individuals with respect to the net-of-tax price. The numerator again captures the equity effect of an increase in the marginal tax rate at the bottom and reflects the difference between the social benefit of a marginal increase in the consumption of the non-working—that is, the basic income guarantee, $-T(0)$—and the fiscal cost of such an increase financed by an increase in $T'(0)$. The denominator is the efficiency effect and reflects the behavioural responses to this change. Given that the government wants to redistribute from high- to low-income individuals, the marginal social welfare weights will be decreasing in income, and since the average marginal social welfare weight is equal to one, the marginal tax rate on the non-working individuals will necessarily be positive. Significantly, the higher the marginal tax rate at the bottom of the income distribution, the greater the social welfare weight of those not working.

A maximin social welfare function will yield a positive marginal tax rate at the bottom even without bunching. The maximin problem for the government is to maximize the utility of the least well-off person, $u\left(c(\underline{w}),y(\underline{w})/\underline{w}\right)$, subject to the government budget constraint and incentive constraints. With additive individual utility functions, Boadway and Jacquet (2008) show that the optimal marginal tax rate at wage rate w when utility is quasilinear in consumption is:

$$\frac{T'\left(y(w)\right)}{1-T'\left(y(w)\right)} = \underbrace{\int_{w}^{\overline{w}} dF(\tilde{w})}_{\text{equity}} \cdot \underbrace{\frac{1+\varepsilon_\ell}{\varepsilon_\ell w f(w)}}_{\text{efficiency}} = \left(1-F(w)\right) \cdot \frac{1+\varepsilon_\ell}{\varepsilon_\ell w f(w)}. \tag{14}$$

This expression is similar to that in (12) except for the equity term. Since no weight is put on the welfare of those with $w > \underline{w}$, there is no loss in social utility on their account when the marginal tax rate applying at w is increased. The equity term reflects only the social value of the increase in government revenue for all individuals between w and $\overline{w}$ when the marginal tax rate on w increases. When the marginal tax rate $T'\left(y(w)\right)$ increases, all those above w, of whom there are $1-F(w)$, pay the same increment in lump-sum tax revenue. Consequently, the marginal tax rate at the bottom of the wage distribution will be positive and given by $T'\left(y(\underline{w})\right) = (1+\varepsilon_\ell)/\left(1+\varepsilon_\ell+\varepsilon_\ell \underline{w} f(\underline{w})\right)$. This implies a 100% marginal tax rate when $\underline{w}=0$ so the lowest-wage individuals never work, and less than 100% when $\underline{w}>0$.

With other social welfare functions, it is generally not possible to obtain further analytical results on the marginal tax rate schedule beyond those at the very bottom and very top of the wage distribution. A partial exception is when the social welfare function is utilitarian. Diamond (1998) shows that when (a) preferences are quasilinear in consumption, (b) the elasticity of labour supply is constant, and (c) the wage distribution is unbounded Pareto above the median, marginal tax rates will follow a U-shaped pattern above the median wage rate. Boadway et al. (2000) show that the Diamond result can also apply when preferences are quasilinear in labour and the social welfare function is utilitarian, but only if the wage distribution is unbounded. Further, as noted by Piketty and Saez (2013, p. 441), it is not possible to obtain analytical results on the size of the transfer at the bottom of the wage distribution. Researchers have generally relied on numerical simulations to say something about the shape of the optimal marginal tax rate schedules and transfers at the bottom under alternate social welfare functions, as we discuss below.

Optimal Nonlinear Taxation: The Perturbation Approach

The above analysis specifies marginal tax rates based on wage rates w as in (11) rather than on income y. Following Saez (2001), the optimal income tax can be restated in terms of income by converting the wage distribution into an income distribution. However, income and the income distribution are endogenous, and that complicates the derivation of optimal income taxes. Saez addresses this issue by deriving optimal marginal tax rates in terms of income by using a perturbation approach. Suppose the optimal income tax is in place, so the income distribution is determined. Then, in the neighbourhood of any income level, increasing the marginal tax rate $T'(y)$ and the intercept b incrementally should not affect social welfare.

Suppose utility is quasilinear in consumption and isoelastic in labour, $\ell = y/w$:

$$u(c, y/w) = c - \frac{(y/w)^{1+1/\varepsilon}}{1+1/\varepsilon}. \tag{15}$$

In the optimum, let the distribution of income be $H(y)$ with density $h(y)$,[3] and the tax function in terms of income is $T(y)$. Maximizing (15) subject to the budget constraint $c = y - T(y)$ yields:

$$y = \left(1 - T'(y)\right)^{\varepsilon} w^{1+\varepsilon}. \tag{16}$$

So, ε is the elasticity of y with respect to $1 - T'(y)$, referred to as the *elasticity of taxable income*.[4] Finally, let $G(y)$ be the average social value in terms of government revenue of giving one unit of income to all persons with income above y:[5]

$$G(y) = \frac{\int_y^{\bar{y}} \frac{W'(\tilde{y})}{\lambda} dH(\tilde{y})}{1 - H(y)}.$$

Since marginal social utility $W'(y)$ is decreasing in y, $G(y)$ will be decreasing in y. Note that $1 - G(y)$ is analogous to the equity term in (12).

Consider now the following perturbation: Increase the marginal tax rate $T'(y)$ by $dT'(y)$ over the interval $y + dy$, holding $T'(y)$ at all other income levels constant. This perturbation will

[3] If income, $y = w\ell(w)$, is increasing in w, we have $F(w) = H(y)$ and $f(w) = h(y)dy/dw$.

[4] In Saez et al. (2012), the elasticity of taxable income refers more generally to when individuals can vary their income by other actions besides changing labour supply. Examples include tax avoidance and evasion, postponing tax liabilities, and changing the form in which income is earned when different forms face different tax rates.

[5] Social welfare continues to be defined as an additive social welfare function that is increasing and concave in individual utilities, whereas utility is now defined as a function of income, and the aggregation is over the income distribution.

have the following effects: Those below income level y will not be affected. For those above y, tax liabilities will increase by a lump-sum amount $dT'(y)$ since all such persons at that level will pay more income taxes on their inframarginal income. Their before-tax income will not change since utility is quasilinear, and the increase in government revenue will be $dR = \big(1 - H(y)\big)dT'(y)$ because there are $1 - H(y)$ individuals with income greater than y. The loss in social welfare for these individuals will be $dW = -G(y)dR$. For those in the interval $y + dy$, of whom there are $h(y)$, the change in income using the elasticity ε is $h(y)dy = -\varepsilon h(y)ydT'(y)/\big(1 - T'(y)\big)$, so tax revenue falls by $dB = T'(y)h(y)dy$, or

$$dB = -\frac{\varepsilon y dT'(y)}{1 - T'(y)}h(y)T'(y).$$

There is no change in utility of persons in this interval since they are only marginally affected by the tax rate change. In an optimum, the sum of these changes is zero—that is, $dR + dW + dB = 0$. Rearranging this we obtain:

$$\frac{T'(y)}{1 - T'(y)} = \frac{\big(1 - H(y)\big)\big(1 - G(y)\big)}{\varepsilon y h(y)} = \frac{\text{equity}}{\text{efficiency}}. \tag{17}$$

The equity term involves the net benefit of transferring tax revenues from the $1 - H(y)$ persons above income level y to the government, where $1 - G(y)$ is the per capita benefit of such a transfer. The efficiency effect involves the elasticity of taxable income multiplied by the total amount of taxable income at income level y, as above.

The advantage of this approach is that the pattern of optimal marginal tax rates is based on y instead of w, reflecting what is done in practice. Many of the properties of (17) follow directly from those discussed above. The optimal marginal tax rate at the bottom and top of the income distribution will both be zero since the equity term in (17) goes to zero at $y = \underline{y}$ and $y = \bar{y}$.[6] The optimal marginal tax rate must be between zero and one in the interior of the income distribution. With a maximin objective, $G(y) = 0$ for all $y > \underline{y}$ and the optimal income tax rule for $y > \underline{y}$ is:

$$\frac{T'(y)}{1 - T'(y)} = \frac{1 - H(y)}{\varepsilon y h(y)}. \tag{18}$$

The optimal marginal income tax rate at the bottom of the income distribution is $T'(\underline{y}) = 1/(1 + \varepsilon \underline{y} h(\underline{y})) > 0$. (Recall that we are assuming ε is identical for all.)

[6] The former follows because, by perturbing b incrementally, the average social value in terms of government revenue of giving everyone with income above $\underline{y}$ is simply one. The latter follows from noting that there is no one with income above $\bar{y}$ when the income distribution is bounded.

We can use this perturbation approach to characterize the optimal marginal income tax rate when it applies over a range of income at the very bottom of the income distribution—for example, when the income tax is piecewise linear. Assume a constant marginal tax rate $T'(y_B)$ applies to incomes below y_B. Let the average income of those in this bottom income bracket be $\bar{y}_B = \int_{\underline{y}}^{y_B} yh(y)dy/H(y_B)$. Now consider a small increase $dT'(y_B) > 0$ at the optimum. The marginal tax increase will affect those earning incomes both above and below y_B. The direct mechanical effect on revenue will be $dR = \left(\int_{\underline{y}}^{y_B} yh(y)dy + \big(1 - H(y_B)\big)\right) dT'(y_B)$. The first term reflects the additional revenue raised on those earning less than y_B. The second term reflects the fact that the marginal tax rates of those above y_B are held constant, but they pay an additional $dT'(y_B)$ in taxes on their inframarginal income. The total welfare cost of collecting this additional revenue is $dW = -g_B \left(\int_{\underline{y}}^{y_B} yh(y)dy\right) dT'(y_B) - G(y_B)\big(1 - H(y_B)\big)dT'(y_B)$ where g_B is the average social welfare weight of those earning less than y_B (weighted by income). Assuming no income effects, the marginal increase in $T'(y_B)$ does not affect the income earned above y_B, but reduces it for those earning less than y_B, and the subsequent revenue cost of this behavioural effect is given by $dB = H(y_B)T'(y_B)d\bar{y}_B$ where $d\bar{y}_B = -\varepsilon_B \bar{y}_B dT'(y_B)/(1 - T'(y_B))$ and ε_B is the average elasticity (weighted by income) of those earning less than y_B with respect to the net of tax price, $1 - T'(y_B)$. Given that this tax change is being considered at the optimum, it must be that $dR + dW + dB = 0$, which yields

$$\frac{T'(y_B)}{1 - T'(y_B)} = \frac{(1 - g_B)y_B^m H(y_B) + (1 - G(y_B))\big(1 - H(y_B)\big)}{\varepsilon_B \bar{y}_B H(y_B)}. \tag{19}$$

The optimal tax rate applying to incomes below a certain threshold takes into account both the equity and efficiency effects in this bottom income bracket in the standard way as given by the first term in the numerator and the denominator of (19), respectively. The optimal tax also reflects equity concerns for those earning more than the threshold as given by the second term in the numerator of (19). Since marginal tax rates are held constant for those earning income above y_B, there are no efficiency concerns arising from their behavioural responses. Rather, they all face a higher total tax liability since the marginal tax rate on income up to y_B is slightly higher. The fewer individuals earning more than y_B, the less revenue the government will raise from them, and the greater their social weight, the more concerned the government is about raising this additional tax revenue from them. In both cases, the optimal marginal tax rate in the bottom income bracket will be lower. Not surprisingly, if no individuals earn more than y_B, so $H(y_B) = 1$, then (19) reduces to the optimal linear progressive income tax rule given by (8).[7]

[7] A similar perturbation exercise could be done for the basic income guarantee b by assuming the total tax liability

Extensive-Margin Labour Supply Approach

In the above subsection, we assumed that labour supply adjusted along the intensive margin such as by varying the number of hours worked. Here we suppose that labour supply adjustments are discrete. They may involve participation, as in Diamond (1980), or job choice, as in Saez (2002), Rothschild and Scheuer (2013), and Ales et al. (2015). In models of extensive-margin labour choice, there are a discrete number of job types. Income in each job is fixed and known to the government—for example, by taxpayer self-reporting. Either disutility in each job or utility if unemployed differs among individuals and is not observed by the government. As in the intensive-margin case, there is a perfectly elastic demand for workers, so anyone who chooses a job is guaranteed to get it at the given income.

We begin with the simple case introduced by Diamond (1980), in which an individual's only choice is whether to participate in the labour force or be voluntarily unemployed. For each individual there is only one job suited to their skills. Government tax-transfer policy consists of a tax contingent on each individual's income. Our analysis follows the assumptions adopted by Saez (2002).

Suppose there are n_i individuals of type i, $i = 0, \cdots, I$. Population is normalized to unity, so $\sum_{i \geqslant 0} n_i = 1$. An individual of type i can take a type-i job that pays y_i and incurs a tax of $T_i \gtreqless 0$, or can choose not to work, earn no income, and obtain a transfer $-T_0$, which is option 0. (Individuals of type 0 are unable to work so necessarily choose option 0.) Utility of work participants is equal to consumption, which for a type-i worker is $c_i = y_i - T_i$. Utility if voluntarily unemployed is $c_0 + \tilde{m}_i = -T_0 + \tilde{m}_i$, where the value of leisure $\tilde{m}_i \in [\underline{m}, \overline{m}]$ is distributed by $\Gamma_i(m_i)$. Thus, individuals differ by their utility of leisure, which is unobservable to the government.

The marginal type-i participant is indifferent between participating or not, so $y_i - T_i = -T_0 + \hat{m}_i$. All those with $m_i < \hat{m}_i$ choose to work. Therefore, the number of type-i participants is $n_i \Gamma_i(\hat{m}_i) = n_i \Gamma_i\big(y_i - T_i + T_0\big) \equiv h_i(\cdot)$, and the number of non-participants is $1 - \sum_{i \geqslant 1} n_i \Gamma_i\big(y_i - T_i + T_0\big) \equiv h_0$. Note that participation is a function of $y_i - T_i + T_0 = c_i - c_0$—that is, the gain in consumption from participating.

of those in the bottom income bracket is given by $T(y) = T'(y_B)y - b$. Without income effects, an increase in b has a mechanical (negative) effect on revenue equal to db times the total population (which has been normalized to unity), and a positive welfare benefit to all those who receive the additional b. Summing these two effects together and setting the sum equal to zero implies that that average welfare weight in the population should equal to one—that is, $1 = \int_{\underline{y}}^{\overline{y}} g(y)h(y)dy$ where $g(y) = (\partial W/\partial y)/\lambda$.

We assume that for social welfare purposes the government gives full weight to the utility that non-participants obtain from leisure, which varies among individuals. This is not an innocuous assumption since it assumes that those who get more benefit from leisure obtain correspondingly more weight in the social welfare function. An alternative approach assumes that individuals have some responsibility for their own preferences and ought not to be rewarded for their preference choices (Roemer, 1998; Fleurbaey & Maniquet, 2011). We return to this issue later when discussing preference heterogeneity.

For now, we assume social utility for participants is $W(c_i) = W(y_i - T_i)$, and for non-participants is $W(c_0 + \tilde{m}_i) = W(-T_0 + \tilde{m}_i)$.

$$\mathscr{W}(\boldsymbol{u}) = \sum_{i \geq 1} h_i(y_i - T_i + T_0) W\big(y_i - T_i\big) + \sum_{i \geqslant 0} n_i \int_{\hat{m}_i}^{\overline{m}} W\big(-T_0 + m_i\big) d\Gamma_i(m_i) \tag{20}$$

where the first term includes the participants, and the second term the non-participants. The government chooses taxes and transfers $T_0, \cdots, T_i, \cdots T_I$ to maximize social welfare subject to its budget constraint:

$$\sum_{i \geq 1} h_i(y_i - T_i + T_0) T_i + \Big(1 - \sum_{i \geq 1} h_i(y_i - T_i + T_0)\Big) T_0 = R.$$

As before, we define *marginal social weights in terms of government revenue* for participants and non-participants, respectively, as

$$g_i \equiv \frac{u_i'}{\lambda} \quad \text{and} \quad g_0 \equiv \frac{1}{h_0} \left(\frac{u_0'}{\lambda} + \sum_{i \geqslant 1} n_i \int_{\hat{m}_i}^{\overline{m}} \frac{u_{i0}'}{\lambda} d\Gamma_i \right)$$

where λ is the Lagrange multiplier on the government budget constraint. The marginal social welfare weight for non-participants is the average marginal social welfare weight in terms of government revenue across all those not working, including both the type 0s who are unable to work and those type-i individuals, $i \geqslant 1$, who choose not to work. We define the *elasticity of participation* for a type-i individual as

$$\eta_i \equiv \frac{(y_i - T_i + T_0) h_i'(y_i - T_i + T_0)}{h_i(y_i - T_i + T_0)} = \frac{(c_i - c_0)}{h_i(c_i - c_0)} \frac{dh_i(c_i - c_0)}{d(c_i - c_0)}.$$

It is the proportionate change in the number of participants $h(\cdot)$ from a proportionate change in the consumption gain from participation, $c_i - c_0$. The size of η_i depends on the distribution of the utility of leisure m_i, of which little is known.

Using these definitions, the first-order conditions for the choice of T_i reduce to:

$$\frac{T_i - T_0}{c_i - c_0} = \frac{\tau(y_i)}{1 - \tau(y_i)} = \frac{1 - g_i}{\eta_i} = \frac{\text{equity}}{\text{efficiency}}, \qquad i \geqslant 1 \tag{21}$$

$$\sum_{i \geqslant 0} h_i g_i = 1 \tag{22}$$

where $\tau(y_i) = (T_i - T_0)/y_i$ is the *participation tax rate* at income level y_i—that is, the increase in tax liability when changing from non-participation to participation. Eqs. (21) and (22) have the following implications: Suppose that non-participants have the highest social weight—say, because they include those unable to work—so $g_0 > g_1 > g_2 > \cdots > g_I > 0$. Then, since the average value of the g_is is unity by (22), there is some i^* such that $g_{i^*} = 1$. For all $i < i^*$, $g_i > 1$, so from (21), $T_i < T_0 < 0$. That is, the participation tax rate is negative. The transfer of $-T_0$ to the type 0s is the basic income guarantee and is smaller than the income transfers to the working poor, $-T_i$ for $i < i^*$.

It is possible, however, that the social weight put on non-participants is high enough that $g_0 > 1$ and $g_i \leq 1$ for all $i \geq 1$ so the participation tax is non-negative everywhere. This is the case if the social welfare function is maximin, and the type 0s are the least well-off. In these cases, $g_i = 0$ for all $i > 0$, so $T_i > T_0$ for all i. The basic income is still $-T_0$. Thus, for a participation subsidy to be optimal, the social welfare weight must be higher on non-participants than participants, but not high enough that $g_i < 1$ for all participants.

If, for example, a high proportion of non-participants have higher skills, it is possible that g_0 is lower than g_i for low-income participants. For example, Christiansen (2015) shows, using a utilitarian social welfare function, that the optimality of a larger transfer to the working poor than to the non-working hinges crucially on a sufficiently high participation elasticity of those individuals with productivity just above the productivity of the working poor. By reducing the transfer to the non-working relative to the working poor, the government increases the incentive of those with slightly higher productivity than the non-working poor to participate in the labour market. The more responsive the participation of these individuals, the greater the potential gain from having them working. At the opposite extreme, if all individuals with productivities greater than type 1 always worked, then g_1 would necessarily be less than one. We return below to the treatment of high-productivity persons with a high preference for leisure.

These results, following Saez (2002), are based on the assumption that individuals differ in both productivity and preferences for leisure. Choné and Laroque (2005) assume instead that, in addition to differing in productivity, individuals differ in their *taste for work*, which affects their utility when working and is modelled as the financial compensation needed to work. All unemployed individuals are equally well-off, and since individuals choose whether to work or not, those who are voluntary unemployed are the worst off. A government with a maximin objective, as Choné and Laroque assume, therefore maximizes the welfare of the voluntary unemployed. They show that in this case, a participation subsidy at the bottom is generally not desirable. That agnosticism applies to the utilitarian case as well, where the utility of participants now counts in social welfare. Since there is no empirical evidence to indicate whether individuals differ in their preferences for work or for leisure, and how such preferences vary with productivity level, the implication is that the case for negative participation taxes at low income levels is weak.

The case for participation subsidies is weakened further when we allow individual choices to include more than participation. Suppose that individuals can choose both whether to participate in the job market and in which occupation. Saez (2002) assumes those of a given productivity level can choose an occupation of the next skill below. Thus, type-i individuals who participate can choose job i or job $i-1$. Labour supplied to occupation i is $h_i(c_i - c_0, c_{i+1} - c_i, c_i - c_{i-1})$. The first term captures the participation margin, while the latter two capture occupational choice margins. The partial elasticity of $h_i(\cdot)$ with respect to $c_i - c_0$ is η_i as before, and the elasticity with respect to $c_i - c_{i-1}$ is denoted as ζ_i. Using a similar approach to the above, optimal taxes satisfy:

$$\frac{T_i - T_{i-1}}{c_i - c_{i-1}} = \frac{1}{\zeta_i h_i} \sum_{j=1}^{I} h_j \left[1 - g_j - \eta_j \frac{T_j - T_0}{c_j - c_0} \right], \qquad i \geqslant 1 \tag{23}$$

where the left-hand side is the tax rate arising from the job choice decision and the tax term on the right-hand side is the participation tax as above. In this case, a participation subsidy for the working poor (i.e., type-1 individuals) is optimal only if the participation elasticity η is large enough relative to job choice elasticity ζ (i.e., the responsiveness of job choice to the difference in after-tax incomes).

To gain some insight as to whether such a participation subsidy is optimal, we evaluate (23) at $i = 1$ and use the normalization of the welfare weights $\sum_{j=0}^{I} g_j h_j = 1$ to obtain:

$$\frac{T_1 - T_0}{c_1 - c_0} = \frac{1}{h_1(\zeta_1 + \eta_1)} \left[(g_0 - 1) h_0 - \sum_{j \geq 2}^{I} h_j \eta_j \frac{T_j - T_0}{c_j - c_0} \right]. \tag{24}$$

Eq. (24) generalizes the insight from the pure extensive-margin model of Christiansen (2015) that the optimality of a negative participation tax at the bottom depends critically on the participation elasticities of those workers with productivities above those of the working poor (types $j > 1$).

Alternatively, suppose, following Jacquet et al. (2013) and Choné and Laroque (2011), that individuals choose whether to participate in a job suitable to their productivity level and how intensively to work. Not surprisingly, Jacquet et al. find that combining an extensive decision into the standard intensive model reduces marginal income tax rates, resulting in less progressivity, and reduces the possibility that a participation subsidy is optimal. They also find that those with zero incomes should be given a different transfer (i.e., basic income) than those with low incomes since the two groups have different characteristics. Choné and Laroque show that work subsidies for low-skilled individuals can be optimal, provided the social welfare weights on workers decrease in both incomes and productivities. They rely on several assumptions on the underlying distribution of work opportunities, including being log-concave, for which there is a lack of empirical support.

Beaudry et al. (2009) consider the optimal tax-transfer system when individuals differ in both their market and their non-market productivities. Market productivities (wage rates) are observable, but non-market productivities are not. They show that it is optimal to provide zero transfers to those not working, to provide wage subsidies for those with wages below some cut-off, and to tax wages above the cut-off.

In general, the more decisions individuals make, the weaker is the case for a participation subsidy at the bottom, and the less progressive is the tax. This is in addition to the case for the participation subsidy being eliminated if non-participants have a lower social weight than participants, or if non-participants have very high social weight, as in the maximin case. A further determinant of participation subsidies arises when participation does not guarantee a job will be landed. We return to this below.

Estimating the Optimal Income Tax Structure

The solution to the optimal income tax problem reflects the trade-off between efficiency and equity. The former depends on the behavioural response of taxpayers to income taxation, and the latter depends on the social weights that the government places on income changes by income group. In this section, we review the methods used in the optimal income tax literature to estimate behavioural effects, to determine social welfare weights, and to compute optimal income tax structures.

Empirical Analysis of Behavioural Responses

The efficiency losses from redistribution depend on individuals' responses to taxation, which can be captured by behavioural elasticities. These elasticities can be estimated empirically. One approach, following the labour economics literature, focuses on estimating labour supply responses to taxes (see Keane, 2011, for a survey of this literature).[8] There is also a growing public economics literature estimating the elasticity of taxable income and participation responses to transfer programs. We briefly review the latter literature.

Estimates of the Elasticity of Taxable Income

The optimal income tax rate formulas are typically derived from models in which taxable income is determined solely from labour supply decisions. In reality, individuals make decisions along many different margins that affect their taxable income, such as the timing of income receipt, the shifting of income between income sources that are taxed at differential rates, undertaking other forms of tax avoidance, or simply not declaring income. Rather than trying to determine the specific margin along which individuals change their taxable income, researchers instead focus on estimating how total taxable income responds to tax changes, which can be captured by the elasticity of taxable income.

Suppose preferences are defined over disposable income c and reported taxable income z and represented by the utility function $u(c,z)$. Saez et al. (2010) define a linearized budget constraint at the chosen point as $c = (1-t)z + E$, where t is the marginal tax rate and E is the individual's virtual income. The individual's maximization problem yields reported taxable income z as a function of the net-of-tax price $1-t$ and virtual income E, and we define the *elasticity of taxable income* as

$$\mathrm{ETI} = \frac{\partial z}{\partial(1-t)}\frac{1-t}{z}. \tag{25}$$

The ETI captures all behavioural responses to changes in t. Without income effects, the ETI also provides a measure of the efficiency costs of income taxation.[9]

Saez et al. (2012) survey the empirical approaches used in estimating the ETI and summarize selective findings. They note (p. 6) that the literature typically ignores income effects on the grounds that there is no general agreement as to their size.[10] Using a share analysis, which exam-

[8] Spadaro et al. (2015, Table 4) summarize the existing empirical estimates of the extensive (participation) and intensive labour supply elasticities for single individuals.

[9] Even without income effects, the ETI may not be a sufficient statistic for the efficiency cost of income taxation if some behavioural responses result in either additional revenue or positive externalities. For example, there may be shifting of income to other tax bases or to future tax bases or the claiming of deductions that have positive social benefits, such as charitable giving, or tax evasion may involve potential financial penalties (Chetty, 2009).

[10] Two exceptions are Gruber and Saez (2002) and Kleven and Schultz (2014), who both estimate very small income effects for reported income using U.S. and Danish panel tax data, respectively.

ines the effect of a tax change on the share of income accruing to the top 1% of income earners while controlling for other factors, the U.S. evidence suggests that the ETI of these top income earners is between 0.58 and 0.82. Veall and Saez (2005) find an estimate of around 0.47 using Canadian data from 1920 to 2000, but a much smaller elasticity when looking only at the share of the top 1% of wage income. Wage income is more subject to third-party reporting and has fewer tax avoidance opportunities relative to capital income.

Using a difference-in-difference approach comparing the incomes of a treatment group that experienced some tax change to the incomes of a control group that did not, the estimated U.S. elasticities vary according to the availability of deductions, the definition of taxable income, and the time period. Sillamaa and Veall (2001), using the 1988 Canadian tax reform, found a much larger ETI for self-employed income (1.3) than for gross income (0.25), as well as much larger responses for high-income earners (> 1) and seniors (0.82). More recently, Milligan and Smart (2015) use variation in provincial income tax rates in Canada to estimate the ETI and find much higher elasticities at the provincial level for incomes in the top 1% than for others.

Focusing on lower-income earners in the U.S., Saez (2010) uses a bunching analysis on the EITC, which is the largest transfer program to low-income families. The transfer increases with income, reaches a plateau, and then is phased out. Its size also increases with the number of children in the family. Saez documents significant bunching of reported income around the income level where the EITC is first maximized. The estimated elasticities from this observed bunching depend on the number of children (0.21 for one child and 0.15 for two or more children) but are largely driven by self-employed income (1.1 and 0.8 for one and two or more children, respectively). Estimated elasticities for the self-employed increase over the period under study, but those for wage earners—who likely have little influence over their income and who are subject to third-party reporting—remain close to zero over the entire period.

Saez et al. (2012) conclude that the ETI for the U.S. likely lies somewhere between 0.12 to 0.4. The midpoint of 0.25 is often used in calibrations, with an assumed higher elasticity for top income earners (Gordon & Cullen, 2011). Studies from other countries find other ranges of estimated ETIs, which is not surprising since the elasticities are dependent on the actual tax system. The key finding is that these elasticities tend to be higher for similar types of groups, such as seniors, high-income earners, and the self-employed. The behavioural margins along which taxable income is responding for these groups likely differ. For example, for seniors, it is likely to be retirement or pension payment decisions; for high-income individuals, access to tax avoidance opportunities;

and for the self-employed, the under-reporting of income and overstating of expenses that are not subject to third-party reporting.

Estimates of the ETI have been used to inform policy debates about taxing the rich (Diamond & Saez, 2011; Veall, 2012). Much less focus has been on the bottom of the income distribution, which would be relevant for basic income recipients. The existing evidence suggests lower taxable incomes are less responsive to tax changes, likely due to third-party reporting and perhaps an inability to adjust real incomes.[11] At the same time, income effects may well be more important for low-income persons who receive transfer income. This would be particularly relevant in assessing a basic income since the size of the transfer would likely be much more important than the change in effective marginal tax rates.

In a recent paper, Bosch et al. (2019) find no evidence of bunching (or, equivalently, estimate a zero elasticity of taxable income) by individuals, including those with self-employment income, in a targeted cash transfer program, despite finding bunching of the same individuals around income tax bracket thresholds in the Netherlands. This suggests that the behavioural responses to clawback rates in a cash transfer program may differ from behavioural responses to statutory marginal tax rates, despite the theoretical prediction that the two responses should be the same.

Estimates of Participation Responses

Kleven (2019) studied the effect of EITC changes in the U.S. on participation rates among single women with varying numbers of children. He considered all federal and state changes to the EITC since the program was introduced in 1975. Participation rates of single women with children showed a pronounced increase around 1993 and closed the gap with childless single women. Subsequently, participation rates for the two groups changed very little. This coincided with a significant increase in the maximum EITC credit, particularly for families with two or more children. To what extent was this increase in participation rates for single women with children due to 1993 EITC reforms? Kleven's empirical analysis takes into account other confounding factors, including major welfare reforms of the early 1990s and the buoyant economy. His result suggests that virtually all changes in participation can be attributed to these confounding factors, and that changes due to the EITC were zero.

[11] In recent fiscal estimates of changes to the federal personal income tax system, the Office of the Parliamentary Budget Officer (2016) simply assumes a constant elasticity of taxable income for those earning less than $90,000 of 0.1 and for those earning more than $200,000 to be 0.38.

The absence of labour supply effects of the EITC could be because of low participation elasticities or lack of information about the program. Recipients are unlikely to know how they are affected by the EITC when they consider labour market decisions because the program is complicated and is only effective over a limited range of low incomes. The provision of information on the EITC to potential EITC recipients by tax preparers was the focus of a field experiment by Chetty and Saez (2013). Their study was motivated by what was at the time a generally accepted finding of strong participation responses to the EITC but marginal- or zero-earning responses. Could providing information about the structure of the EITC to EITC-eligible tax filers affect their earning responses? The authors find no effect from providing information on subsequent earnings of potential EITC recipients, on average. They did, however, find evidence that the earning responses to the information provided depended on the tax professional, which suggests that administrators of transfer programs can have real effects on how individuals respond to these programs.

Finally, Card and Hyslop (2005) consider the Self-Sufficiency Project (SSP) in Canada, which provided temporary, but generous, earning subsidies to social assistance recipients who were working full-time. The SSP was shown to have significant positive (negative) effects on employment (welfare) participation in the first 18 months of the three-year program. But this effect faded with time and completed vanished 18 months after the end of the program.

Social Preferences

The structure of the optimal income tax system is sensitive to the assumed social welfare function. This is neither surprising nor particularly concerning when one views the role of optimal tax theory as a guide to designing tax systems for a given normative perspective. But when considering tax policies for actual economies, the question becomes: What is society's social welfare function? One approach is to suppose that the government's redistributive preferences reflect the collective preferences of citizens. One immediately faces the voting paradox problem of Arrow (1951), who showed that it is generally impossible to aggregate individual preference orderings into a transitive social ordering. To avoid this, assumptions are typically made to ensure that preferences over redistribution are single-peaked. This can be achieved, for example, if the government is restricted to a linear progressive income tax (Gans & Smart, 1996) or if government decision-making is entrusted to a single authority, the "citizen-candidate" (Osborne & Slivinski, 1996).

Another approach is to assume that redistribution is motivated by the altruism of the rich (Hochman & Rodgers, 1969). In this case, redistribution is based on efficiency, but because voluntary transfers of the rich are constrained by free-riding behaviour, the government redistributes on their behalf. The drawback of this approach is that only the utility of the rich counts, albeit utility that includes the altruistic benefit the rich get from the well-being of the poor. The altruistic approach leads to less redistribution than social welfare maximization would propose.

Two main approaches have been used to try to uncover the relevant objectives of government in practice. The first is the *optimal inverse approach*, which determines the implied social welfare weights from the actual tax-transfer system chosen by the government. The underlying assumption is that the government uses a social welfare function in choosing actual tax policy. The second is the *stated-preference approach*, which surveys individuals for their views on different social objectives. How one then aggregates these individual preferences into a social objective that could be used for tax policy evaluations remains an open question.

Optimal Inverse Approach

In this approach, the government's social preferences are estimated from the actual tax/transfer system. The social preferences that would generate the existing tax system are backed out, given assumptions about individual preferences, production technology, and skill distributions. This approach was first used in the context of optimal commodity taxation by Christiansen and Jansen (1978) using Norwegian data, and extended to optimal income taxation by Bourguignon and Spadaro (2012) using French data.

To understand this approach, suppose the utility function is quasilinear in consumption with isoelastic labour supply, as given by (15). The optimal tax rule (12) can then be rewritten as:[12]

$$\frac{T'(w)}{1-T'(w)} = \left(1+\frac{1}{\varepsilon}\right)\left(1-S(w)\right)\frac{1-F(w)}{wf(w)} \tag{26}$$

where

$$S(w) = \frac{\int_w^{\overline{w}} \left(W'(\tilde{w})/\lambda\right) f(\tilde{w}) d\tilde{w}}{1-F(w)}. \tag{27}$$

Recall that when $\mu'(c)=1$, $W'(w)/\lambda$ is the marginal social welfare weight of an individual with ability w—that is, the marginal increase in social welfare in terms of public funds from giving an incremental amount of consumption to an individual of ability w. The expression $S(w)$ represents the average marginal social welfare weight of all individuals with ability greater than w, where

[12] This is the so-called ABC optimal income tax rule derived for these preferences by Diamond (1998), $T'(w)/(1-T'(w)) = A(w)B(w)C(w)$, where $A(w) = 1+1/\varepsilon$, $B(w) = 1-S(w)$ and $C(w) = (1-F(w))/wf(w))$.

$S(\underline{w}) = 1$ since there are no income effects. Note that $1 - S(w)$ is the equity term that appeared in the optimal income tax expression (12) earlier.

Rearranging (26), $S(w)$ can be expressed in terms of the marginal tax rate, the ability distribution, and the elasticity of labour supply, or, equivalently, the elasticity of taxable income:

$$S(w) = 1 - \frac{T'(w)}{1 - T'(w)} \frac{\varepsilon}{1 + \varepsilon} \frac{wf(w)}{1 - F(w)}. \tag{28}$$

To obtain $S(w)$ from (28), Bourguignon and Spadaro (2012) use an estimated value for ε and observed marginal tax rates. Eq. (16) gives income in terms of ε, $T'(y)$ and w, and can be inverted to solve for w as a function of y. This can be used to determine the wage distribution $F(w)$. The marginal social weights can then be obtained by noting that by (27)

$$\frac{dS(w)}{dw} = -\frac{W'(w)}{\lambda} \frac{f(w)}{(1 - F(w))}.$$

One can instead use (17) for the optimal tax rate from the perturbation approach. Rearranging this gives an expression for the average marginal social welfare weight of all individuals with labour earnings greater than y as a function of the observed tax schedule and the income distribution (see Lockwood and Weinzierl, 2016):

$$G(y) = 1 - \frac{T'(y)}{1 - T'(y)} \frac{\varepsilon y h(y)}{1 - H(y)}. \tag{29}$$

Differentiating this and assuming that $T'(y)$ is constant around any income level, so the income tax is piecewise linear as in practice, we obtain:

$$g(y) = 1 + \frac{T'(y)}{1 - T'(y)} \varepsilon \left(1 + \frac{yh'(y)}{h(y)}\right) \tag{30}$$

where $g(y) = W'(y)/\lambda$ is the marginal social welfare weight. The last term is the elasticity of the density of the income distribution at income y, which can be determined empirically (Bastani & Lundberg, 2017). Using the actual tax system, observed income distribution, and an estimated elasticity of taxable income, a marginal social welfare weight can be obtained at each income y.

The above captures labour supply decisions only on the intensive margin, but the inversion approach can be extended to the case where individuals make discrete participation and job choice decisions, as in Saez (2002), discussed above. The optimal tax rule is now given by (23), which involves elasticities of participation η_i and job choice ζ_i. The marginal social welfare weight for $i = I$ can be derived directly from (23) as a function of observed income levels, taxes, and elasticities and is given by

$$g_I = 1 - \eta_I \frac{T_I - T_0}{c_I - c_0} - \zeta_I \frac{T_I - T_{I-1}}{c_I - c_{I-1}}. \tag{31}$$

With g_I in hand, one can solve recursively for the marginal social welfare weight for the $i = 1, \ldots, I - 1$ groups using (23). Finally, g_0 can be obtained given the normalization of the marginal social welfare weights in (22) (Bourguignon & Spadaro, 2012).

Two key issues that papers using this approach have studied are (a) whether marginal social welfare weights estimated from actual tax systems imply a social welfare function that is Paretian (i.e., exhibits non-negative marginal social welfare weights) and (b) whether the government puts less weight on higher-income persons (i.e., marginal social welfare weights decreasing in income).

Using the intensive-margin model, Bourguignon and Spadaro (2012) find support for an implied social welfare function in the 1995 French tax system that is increasing and concave in incomes. But this result is quite sensitive to the assumed labour supply elasticity and depends on which aspects of the French tax-transfer system are included in the tax rates used.

Bastani and Lundberg (2017) do a similar exercise using the Swedish tax system but also incorporate a participation decision and assume, based on existing evidence, that participation elasticities are decreasing in income. They find that the implied social weights are the highest on middle-income workers and lower for low- and high-income workers. This is consistent with a government objective that maximizes the welfare of the median voter, where the highest weight would be on the median income person and lower weights on those with higher and lower incomes. Further, weights on non-workers are quite high relative to low-income workers, and welfare weights at the top of the income distribution are negative for some years. Jacobs et al. (2017) use a framework similar to Bastani and Lundberg to uncover the political weights across the income distribution of political parties in the Netherlands based on the parties' policy platforms. All parties give higher political weight to low-income groups relative to high-income groups, but the weights increase from low-income to middle-income groups and can be close to zero at the top of the income distribution. This latter result disappears if one allows for different beliefs about behavioural responses by political parties.

Lockwood and Weinzierl (2016) focus only on individuals with positive income. They find the implied marginal welfare weights to be high and relatively constant over the bottom half of the income distribution and then fall somewhat, but still remain positive and fairly high at the very top. They also document a large shift up in the implied social marginal welfare weight for the top income earners after the 1986 U.S. tax reform, which reflects the substantial reduction in the top marginal tax rate. Varying the assumed elasticity of taxable income does not affect the implied marginal social welfare weights for those at the bottom but, not surprisingly, reduces it for those at the top.

Bargain et al. (2014a, 2014b) examine the implied welfare weights across several European countries and the U.S. They first obtain comparable extensive and intensive labour supply elasticities by estimating them for each country using the same empirical approach and focusing only on single individuals without children. They find the marginal social welfare weights to be mostly positive across all countries and income groups. The relative size of the weights varied across countries, with much higher weights on the non-working group in countries that provide substantial transfers at the bottom of the income distribution and very small weights on low-income workers in countries with large social assistance programs for the working poor. Bargain et al. (2014b) suggest the potential non-concavity of the implied social welfare function could have arisen from governments being uninformed about behavioural responses of such redistribution programs, but it still remains that some countries have a strong preference for redistributing income to non-working individuals.

Using the implied welfare weights, Bargain et al. (2014a) estimate an index of social inequality aversion for each country and find that countries can be roughly separated into three groups according to their level of inequality aversion. To do this, they assume that the social welfare function depends only on group i's consumption and takes the form $W(c_i) = c_i^{1-\gamma}/(1-\gamma)$, such that $g_i = W'(c_i)/\lambda = 1/(\lambda c_i^{\gamma})$ so a higher γ reflects a greater taste for redistributing income from high- to low-income individuals. Given the estimated values for the g_i group, they can estimate this inequality aversion measure for each country. Their estimated inequality aversion measure ranges from 0 to 3 across the various countries they consider, but for the majority of countries, they estimate a $\gamma \in (0,1)$. Countries in Southern/Eastern Europe and the U.S. have low indices of social inequality aversion (close to zero) reflecting a low social marginal weight on the non-working and a higher than average social marginal welfare weight on the working poor. The Scandinavian countries and Belgium have very high indices of social inequality aversion (greater than one) reflecting a very high marginal social welfare weight on the non-working. Finally, the U.K., Ireland, and Continental Europe have indices in between these two other groups of countries.

Spadaro et al. (2015) do a similar cross-country comparison (albeit assuming a common baseline for both the participation and labour elasticities and then adjusting the baseline elasticities for each country by a measure of the country's relative labour market flexibility) and compare the implied welfare weights on non-working singles—that is, g_0—relative to the weight on the working poor, g_1, and finally four other higher income groups, g_2, g_3, g_4, and g_5. They find that across all countries, the implied social marginal welfare weights on the four higher-income groups are relatively constant. The key differences across countries are the implied weights on the non-working and the working poor. In countries with generous social assistance programs that generate high effective marginal tax rates (such as the Nordic countries and Germany), the implied social welfare weights are quite high for the non-working and relatively low for the working poor. This in contrast to countries with less generous transfer programs, which have slightly higher (or the same) implied social welfare weights for the non-working as the rest of the income groups. Blundell et al. (2009) conduct a similar analysis on single mothers in the U.K. and Germany. They find that the implied social welfare weights are high on non-working mothers and lower and relatively constant for working mothers.

While there is far from a single shape of social welfare weights implied by existing tax-benefit systems across all countries, the existing literature suggests that they follow a general pattern, with weights associated with the non-working being possibly the highest, and weights falling in the working poor or bottom-income-earner groups and then remaining relatively constant or falling in income. The decrease is not generally monotonic and varies across countries. The highest-income groups have the lowest weights, but they may still be quite high (as in the U.S.) or negative (as in Sweden). These implied social marginal welfare weights generally reveal social preferences with a strong preference for redistributing toward those at the very bottom of the income distribution for some groups of countries. These results support the provision of some form of basic income to non-working poor, but at the same time simply reflect the degree to which these countries are already transferring income to those not working.

Stated Preference Approach

A more direct way to solicit social preferences is to ask people for their redistribution preferences, as is done in the General Social Survey in the U.S., the European Social Survey in several European countries, and the World Values Survey in several other countries. Such surveys ask individuals for their views about the role of government in redistributing income, either to ensure everyone (or just the poor) is provided for or to reduce income inequality. Individuals are asked to respond along a numerical scale that corresponds to the level by which they agree or disagree with a statement that it is the role of government to do so. Researchers then try to explain how these responses depend on observable characteristics, such as age, marital status, gender, employment status, education, and race. Alesina and Giuliano (2011) survey the empirical results for the U.S., and Olivera (2015) summarizes results from various European countries.

The impetus for this approach was as an empirical test of the median-voter model of income redistribution developed by Meltzer and Richard (1981). This model predicts that the more unequal the underlying distribution of the individual productivities (as measured by the difference in the mean and median of the distribution), the more income redistribution there will be. To see this, consider the linear progressive income tax system discussed above and assume there is a constant labour supply elasticity with no income effects. An individual with productivity w has the indirect utility $v(t,b;w)$, as given by (4). With no income effects, labour supply will be strictly decreasing in the tax rate, and higher-w individuals will provide more labour and earn higher incomes.[13]

Now suppose individuals could choose their preferred tax rates, recognizing that the lump-sum transfer b is determined from the government's budget constraint in (5), with a zero revenue requirement ($R = 0$). The individual's preferred tax rate solves

$$\max_t \; v(t,b;w) \text{ s.t. } b = t\int_{\underline{w}}^{\overline{w}} w\ell(t;w)dw$$

which yields $t(w)$. Individuals' preferences are single-peaked, and preferred tax rates will be decreasing in w.[14] Applying the median-voter theorem, the equilibrium tax rate, t^m, will be the preferred tax rate of the individual with the median productivity or, equivalently, with median income, which we denote by y^m:

$$\frac{t^m}{1-t^m} = \frac{1-\frac{y^m}{\mathbb{E}(y)}}{\varepsilon}. \tag{32}$$

The numerator reflects the difference between the average income and the median income. It is positive when the median income is below the mean and will be larger the higher the mean income is relative to the median income. The denominator is the elasticity of labour supply or, equivalently, the elasticity of taxable income, which would be weighted by income if it were not assumed to be the same for all individuals.[15] The median-voter model then predicts that the size of the government (in

[13] Assuming $u(c,\ell) = c - h(\ell)$ with $h', h'' > 0$, the solution to the individual maximization problem yields $\ell(t,b;w)$, where $\partial\ell(t,b;w)/\partial t = -w/h''(\ell) < 0$, $\partial\ell(t,b;w)/\partial b = 0$, and $\partial\ell(t,b;w)/\partial w = (1-t)/h''(\ell) > 0$. Given that there are no income effects $\partial v/\partial t = -w\ell$ and $\partial v/\partial b = 1$.

[14] The first-order condition is $\partial v/\partial t + (\partial v/\partial b)db/dt = 0$, where $db/dt = \int_{\underline{w}}^{\overline{w}} w\ell(t;w)dw + t\int_{\underline{w}}^{\overline{w}} w(\partial\ell(t;w)/\partial t)dw = \mathbb{E}(y) - \frac{t}{1-t}\int_{\underline{w}}^{\overline{w}} w\ell\varepsilon dw$, which is the same for all individuals. The first-order condition can be solved implicitly for $t(w)$, and totally differentiating yields $t'(w) = (-\ell - w\partial\ell/\partial w)/(1/(1-t)^2 - (w\ell/\mathbb{E}(y)^2)\int_{\underline{w}}^{\overline{w}} w(\partial\ell/\partial t)dw < 0$.

[15] In defining $\mathbb{E}(y) = \int_{\underline{w}}^{\overline{w}} w\ell(t^*;w)dw$ and $y^m = w^m\ell(t^*;w^m)$, let $\varepsilon = (\partial(w\ell)/\partial(1-t))((1-t)/(w\ell))$. It then follows from using the expressions for $\partial\ell/\partial w$ and$\partial\ell/\partial t$ that $\varepsilon_\ell = (\partial\ell/\partial w)(w/\ell) = (\partial\ell/\partial w)(w/\ell)((1-t)/(1-t)) = -(\partial\ell/\partial(1-t))((1-t)/\ell)(w/w) = \varepsilon$.

terms of amount of income redistribution) will depend positively on the degree of income inequality.

Existing evidence provides limited support for the median-voter model of income redistribution. This could be due to a failure of the median-voter rule in aggregating social preferences rather than individual preferences for redistribution being unaffected by income inequality. Alesina and Giuliano (2011) have shown that at least in some countries, individual preferences are affected by observed inequality. In the U.S., however, evidence from the General Social Survey shows that individuals' stated preferences for redistribution have not been affected by changes in income inequality, even by those with incomes below the mean (Kuziemko et al., 2015). To explore whether this finding is a result of a lack of information about the actual amount of income inequality, Kuziemko et al. use a randomized experimental information survey treatment and find that while stated concern for income inequality increases with information, it does not increase support for government redistributive policies, in large part due to mistrust in the government.

Researchers have also examined how individual perceptions about either immigration or mobility within the income distribution affect stated preferences for redistribution. In a cross-country analysis, Alesina et al. (2019) find that individuals greatly overestimate the size of immigrant populations and misperceive the differences in the characteristics of this population relative to their own characteristics. They find that just asking people about immigration negatively influences stated preferences for redistribution. Alesina et al. (2018) document individual beliefs about intergenerational mobility in three European countries (France, Italy, and Sweden), the U.K., and the U.S. using survey data and find that in the U.S., respondents overestimate the degree of actual intergenerational mobility, whereas respondents in other countries underestimate it. In all countries, they find a negative correlation between beliefs about the degree of intergenerational mobility and preference for redistribution: The greater the optimism about mobility, the lower the desired redistribution, with the magnitude of the correlation depending on the form of government redistribution. Specifically, a stronger correlation was observed for public expenditure policies (such as education and health care) than for income tax and transfer programs. Using a randomized experimental treatment that attempts to manipulate beliefs, they show that the treatment results in less optimistic beliefs about intergenerational mobility as intended, and that it increases support for redistribution for individuals who self-identified as liberal or very liberal with regard to economic issues, but it does not change support for those who reported being conservative or very conservative with respect to economic issues.

Two other interesting lines of investigation relevant to basic income explore whether preferences for redistribution are affected by individual perceptions about (a) whether income (high or low) is largely determined by factors outside or under the control of individuals (i.e., "luck" versus "effort,") and (b) how deserving potential income transfer recipients are. Examining the perceived determinants of income, findings suggest that the more individuals perceive income being determined by effort, the lower the preference for redistribution (Fong, 2001). This also holds for the perceived relative deservedness of potential income transfer recipients. Those unable to work or unable to find work are perceived to be more deserving of an income transfer than those able to work or not looking for work (Saez & Stantcheva, 2016, Online Appendix C). These findings support the Fleurbaey and Maniquet (2011) normative argument that there can be a set of personal characteristics that determine who is the most deserving of transfers separately from income.

Other survey findings also suggest that social preferences are not well represented by the standard social welfare functions. In the optimal income tax approach, social welfare depends only on final outcomes and not on the process by which they are obtained or on initial positions. Using a stated-preference experiment, Tarroux (2019) shows that individuals' rankings of different allocations of pre- and post-tax income in terms of fairness depend on the information provided to individuals about the tax function used to obtain the post-tax income allocations from the pre-tax allocations. Survey respondents were shown to place value on the progressivity of average tax rates independent of their impact on the final distribution of income. Weinzierl (2014) also presents survey evidence that shows that when placed in the position of a social planner, some individuals will choose a tax system satisfying the principle of equal sacrifice such that everyone gives up the same amount of utility when paying income taxes over a tax system consistent with a utilitarian social welfare function. Of course, one can always question (a) the representativeness of the sample respondents, (b) the wording of the survey instrument and whether it actually captures preferences for specific normative principles, and (c) whether individuals are making selections for reasons unrelated to their support for specific normative principles.

Numerical Simulations

Numerical simulations aim to characterize the pattern of marginal and average optimal income tax rates for reasonable assumptions about parameter values and functional forms. Most simulations focus on nonlinear optimal income taxation in the intensive-margin model, whose solution for optimal marginal tax rates is given by (11). As this expression indicates, three key determinants of the structure of optimal tax rates are the distribution of wages or productivities, household preferences over consumption and leisure, and the social welfare function (Tuomala, 2016). We discuss them in turn, bearing in mind the desire to make empirically relevant choices for each.

Distribution of Wages

The choice of wage distribution will determine pre-tax inequality, and the optimal taxation schedule will take this inequality into account. Three distributions have been prominent in simulations. In the *lognormal distribution*, used by Mirrlees (1971), $\log(w)$ follows a normal distribution. This distribution is single-peaked and skewed to the right. It fits empirical distributions of the majority of individuals well, but the fit breaks down in the top and bottom tails of the distribution.

The *Pareto distribution*, used by Saez (2001), takes the form:

$$F(w) = 1 - \left(\frac{\underline{w}}{w}\right)^{\alpha}$$

where $\underline{w}$ is the minimum value of w. The density $f(w)$ declines at a fixed proportional rate. The Pareto distribution fits the upper tail of wages fairly well, but the rest of the distributions poorly. It peaks at the lowest wage level, which is counterfactual.

Tuomala (2016) argues that the *Champernowne distribution* (Champernowne, 1952) matches the data the best since it approximates the Pareto distribution at the upper tail of wages, while having an interior maximum. The cumulative distribution and density functions are:

$$F(w) = 1 - \frac{m^{\theta}}{(m^{\theta} + w^{\theta})}; \qquad f(w) = \frac{m^{\theta} w^{\theta-1}}{(m^{\theta} + w^{\theta})^2} \tag{33}$$

where the scale parameter m is the median of the distribution, while the shape parameter θ can be interpreted as the inverse of the Gini coefficient of the distribution (i.e., Gini$= 1/\theta$). When θ is greater than one, the Champernowne distribution is unimodal, and as θ increases, the dispersion (or spread) of the distribution decreases. Therefore, θ represents inequality in the skill distribution before taxes or transfers. Figure 5 compares the density function of the Champernowne (f_c) distribution (under various shape parameter values) to that of the lognormal (f_l) distribution.

One implementation issue to be overcome is that one cannot infer the distribution of wages from the distribution of incomes since income is the product of labour supply and the wage rate, $y = w\ell$. The existing tax system affects how much labour an individual supplies, and the latter will be conditional on individuals' preferences between leisure and consumption, as well as the tax system. Changing the assumed utility function or tax parameters will change the resulting income distribution. Saez (2001) proposes a method to back out a distribution of productivities. First, he makes assumptions about the utility of individuals, and then he calibrates the distribution of productivities so that the actual tax schedule yields an income distribution that matches empirical measurements. Tuomala (2016) points out two substantial problems with this methodology. First, Saez assumes that the elasticity of labour supply is constant and the same for all workers. Second, he assumes a linear tax schedule when inferring the distribution of productivities from the earnings distribution. This is clearly incorrect given the structure of real tax systems but is required for structural identification of the model.

An alternative method is used in Kanbur and Tuomala (1994). They accept the nonlinear tax schedule present in most existing tax systems and select utility functions so that there is appropriate variation in labour supply elasticities. Since it is not possible to back out the distribution of productivities, they simply select a skill distribution, which, given their model assumptions, creates an earnings distribution that matches the empirical distribution of earnings.

Individual Preferences

Individual preferences between consumption and leisure are defined by the utility function chosen for the numerical tax simulations. The optimal tax schedule will be affected by preferences, since they determine how individuals respond to taxation.

Utility functions are chosen to be tractable. Two common forms are used. One is *quasilinear in consumption*, $u\big(c - h(\ell)\big)$, where c is consumption, $h(\ell)$ is the disutility of labour with $h'(\ell) > 0$, $h''(\ell) > 0$, and $u(\cdot)$ is concave. With this quasilinear utility function, labour supply depends only on the after-tax wage rate, and the uncompensated labour supply elasticity is the same as the compensated elasticity. Sometimes the disutility of labour function is assumed to be constant elasticity, in which case labour supply elasticities are as well.

The second case is the *constant elasticity of substitution* utility function:

$$u(c,\ell) = \Big(c^{1-\sigma} + (1-\ell)^{1-\sigma}\Big)^{\frac{1}{1-\sigma}}. \tag{34}$$

This utility function is homothetic so gives rise to unit income elasticities of demand for consumption and leisure. The uncompensated elasticity of labour supply will now reflect both compensated changes and income effects.

A less commonly used utility function is *quadratic* in both consumption and leisure.

Social Welfare Function

Social welfare functions are typically assumed to be additive and of the general form:

$$\mathscr{W}(u_1, \cdots, u_n) = \sum_{i=1}^{n} W(u_i)$$

where u_i is the utility of a type-i person, and $W(u_i)$ is a social utility function with $W'(u_i) > 0 \geqslant W''(u_i)$. The same social utility function applies to all individuals, and u_i will typically be an indirect utility function that depends on after-tax price and income variables. The social utility function usually takes one of two forms. One is the constant elasticity form:

$$W(u_i) = \frac{u_i^{1-\gamma}}{1-\gamma}, \quad \text{where} \quad -\frac{u_i W''(u_i)}{W'(u_i)} = \gamma \tag{35}$$

so γ is the elasticity of marginal social utility, or the *coefficient of relative aversion to inequality*.

The second form is:

$$W(u_i) = -\frac{e^{-\beta u_i}}{\beta}, \quad \text{where} -\frac{W''(u_i)}{W'(u_i)} = \beta. \tag{36}$$

In this case, β is the *coefficient of absolute aversion to inequality*. The marginal social value of an individual's utility declines at a proportional rate throughout the income distribution. This form has been criticized for this proportional decline, since it applies at both the high end and the low end of utility values.

Notice that in both cases, as the coefficients γ or β go to zero, the social welfare function takes the utilitarian form. If the coefficients go to infinity, the social welfare function becomes maximin.

Results From Numerical Simulations

We report on a representative set of results found in Mirrlees (1971), Saez (2001), and Tuomala (2016). While most numerical simulations focus on marginal tax rates throughout the distribution, average tax rates are what show the extent of redistribution and reflect the size of the basic income. We are particularly interested in marginal tax rates for those receiving basic income and how various modelling choices affect these tax or clawback rates.

Mirrlees (1971) assumes a lognormal distribution of wage rates and an additive social welfare function, with social utilities given by (36). For preferences, he assumes a utility function with a unit elasticity of substitution between consumption c and leisure $1-\ell$ (i.e., Cobb-Douglas utility):

$$u(c,\ell) = \log c + \log(1-\ell). \tag{37}$$

Some results of Mirrlees' numerical simulations are shown in Table 1, where $F(w)$ denotes the position in the wage distribution, ATR is the average tax rate, MTR is marginal tax rate, and X/Z is the fraction of national income available for redistribution. If $X/Z < 1$, then the government must spend some revenue on other purposes, while if $X/Z > 1$, government has revenue beyond taxation available for redistribution. Mirrlees reports both the $\beta = 0$ (utilitarian) case and the $\beta = 1$ (positive aversion to inequality) case. The revenue required for redistribution differs slightly in the two cases, but the difference is small enough that one can see how differences in the inequality aversion β affect the results. For both cases, there is a large, negative, average tax rate at the bottom consistent with a basic income. In addition, marginal tax rates for the lowest-skilled individuals are lower than for those at the fifth percentile of skill, reflecting an increasing marginal tax rate at the very bottom (though it begins to fall quickly thereafter). As inequality aversion (β) increases, average tax rates increase substantially at the very bottom, while marginal tax rates continue to increase at the bottom, though only slightly.

Table 1

Mirrlees (1971) Numerical Simulation Results

	$\beta = 0$ (X/Z = 0.93)		$\beta = 1$ (X/Z = 0.98)	
$F(w)$	ATR (%)	MTR (%)	ATR (%)	MTR (%)
0		23		30
0.05	−34	26	−66	34
0.1	−5	24	−34	32
0.2	9	21	7	28

Saez (2001) uses a Pareto distribution to represent ability. Two alternative utility functions are used, one without income effects on labour supply (quasilinear in consumption) and one with income effects:

$$u = \log\left(c - \frac{\ell^{1+k}}{1+k}\right); \qquad u = \log(c) - \log\left(\frac{\ell^{1+k}}{1+k}\right). \tag{38}$$

All individuals thus have the same labour supply elasticity. He presents results for both the Rawlsian and utilitarian cases, though we focus on the latter. He also simulates results for both the linear progressive income tax and the nonlinear income tax, though we report only the latter.

Table 2 reports Saez's (2001) findings for the level of guaranteed income as a fraction of average income for elasticities of labour supply of 0.25 and 0.5, with and without income effects on labour supply. The addition of an income effect in individual preferences increases the size of the guaranteed income substantially. As well, basic income falls as the elasticity between consumption and leisure increases.

Table 2

Saez (2001) Guaranteed Income as a Fraction of Average Income (Utilitarian Case)

	BI/avg.	BI/avg.
Elasticity of labour supply:	0.25	0.5
No income effect:	0.33	0.21
Income effect:	0.4	0.31

Tuomala (2016) uses instead a Champernowne distribution for the distribution of wages. Following the empirical findings of Riihelä et al. (2013), Tuomala chooses various values for θ (pre-tax inequality) in (33) between 1.5 and 3.3. The higher the value of θ, the lower the level of pre-tax inequality. Tuomala compares simulation results for three different utility functions. The first two are special cases of the constant elasticity of substitution utility function (34) above, which he writes in the following equivalent form:

$$u(c,\ell) = \left(c^{-a} + (1-\ell)^{-a}\right)^{-\frac{1}{a}}. \tag{39}$$

The elasticity of substitution between consumption and leisure is given by:

$$\delta = \frac{1}{1+a}.$$

In Tuomala's first case, denoted u_1, the elasticity of substitution δ is unity, which gives the standard log-log utility (37) used in Mirrlees (1971). In the second case, u_2, $a = 1$, so $\delta = 0.5$, and the utility function is:

$$u_2(c,\ell) = -\frac{1}{c} - \frac{1}{1-\ell}. \tag{40}$$

To justify this selection, Tuomala refers to Stern (1976), who estimates δ to be 0.408. His third utility function is quadratic in consumption, with a bliss point:

$$u_3(c,\ell) = (c-1) - a(c-1)^2 - (1-\ell)^{-1}, \qquad \ell \in (0,1). \tag{41}$$

Table 3 shows average and marginal tax income rates for persons located at the 10th percentile of skill ($F(w) = 0.1$) under a utilitarian scheme. With each utility function, and for different levels of pre-tax inequality, those at the bottom have a negative average tax rate, and a positive and substantial marginal tax rate. Comparing the results of u_1 and u_2 for equal pre-tax inequality, we see that as the elasticity of leisure and consumption decreases, the average and marginal tax rates increase at the bottom, meaning the results of the Mirrlees (1971) case hold, despite the change in the wage distribution used. Similarly, comparing the results for the quadratic utility u_3 with the results of the more common u_2, the quadratic utility function results in lower (though still substantial) average tax rates and marginal tax rates.

Table 3

Utilitarian Tax Schedule Results of Tuomala (2016)

Utility	θ	ATR (%)	MTR (%)
u_1	3.3	−35	38
u_2	3.3	−50	58
	2.5	−133	62
u_3	2.5	−95	11

The results are even more extreme in the maximin case with quadratic utility functions, as seen in Table 4. Here we see that despite the same pre-tax inequality, the average and marginal tax rates at the bottom are much higher, reflecting the greater degree of tax progressivity under maximin.

Table 5 reports Tuomala's results for basic income as a fraction of average income under different social welfare functions for the case with an elasticity of consumption and leisure of 0.5 (u_2). The maximin social objective supports the largest ratio of a basic income to average income, though the utilitarian case results in a large basic income as well. Changing pre-tax inequality does not substantially change the level of the basic income.

Table 4

Maximin Tax Schedule Results of Tuomala (2016)

Utility	θ	ATR (%)	MTR (%)	x_0
u_{12}	2.5	-139	88	0.137

Table 5

Basic Income Results of Tuomala (2016)

	Maximin	Utilitarian
θ	BI/avg.	BI/avg.
2.5	0.68	0.51
3.3	0.68	0.46

Figure 6 shows the shape of the average tax schedules in the maximin (a) and utilitarian (b) cases, for various shape parameters of the Champernowne distribution. In both cases, there is a large negative average tax rate (basic income) at the bottom, with the level of the basic income decreasing as pre-tax inequality decreases. The average tax becomes less negative quickly at the very bottom, rising slowly in the middle, before increasing rapidly at the very top.

The most striking result of these simulations is the lack of a universal shape for marginal tax rate schedules. Different preferences, skill distributions, and social welfare functions lead to very different marginal tax schedules. Figure 7(a) provides a schedule of marginal tax rates for various levels of pre-tax inequality. With high pre-tax inequality ($\theta = 2.0$), the marginal tax schedule increases steadily from the bottom, before dropping at the very top. The low pre-tax inequality case ($\theta = 3.3$) has a mild U-shape. Figure 7(b) compares the marginal tax rates for the utilitarian case to the maximin case. In the utilitarian case, marginal tax rates increase at the bottom quickly, then level off at around the 20th percentile and increase slowly. At the very top, they drop again. In the maximin case, marginal tax rates decrease slowly over the distribution, though fall a bit more rapidly at the very top.

This lack of a definitive structure for the optimal tax system also holds with the piecewise linear tax system. Bastani et al. (2019) calculate the welfare gains from moving from a uniform linear tax system to a two-bracket linear tax system, to a four-bracket linear tax system, and to a fully nonlinear optimal tax system. Their numerical simulations show that under a maximin objective, the income transfer at the bottom also goes up as the number of income brackets increases, but the opposite happens when the government maximizes the sum of the log of individual utilities (Tables 1 and 3).

The numerical simulations of Slemrod et al. (1994) contrast with those of Bastani et al. (2019). They compare the structure of a two-bracket linear income tax with standard linear progressive income tax yielding the same tax revenue using simulation techniques. They assume preferences have constant elasticity of substitution in consumption and leisure with varying elasticities of substitution, and the wage rate follows a lognormal distribution. Two different additive social welfare functions are considered, one utilitarian and another with some finite inequality aversion. For all combinations of parameter values, they find that in the two-bracket case, the optimal marginal tax rate is higher at the bottom than at the top, and the basic income—the transfer received by anyone earning zero income—is higher in the two-bracket case than in the linear progressive case. The intuition is that lowering the tax rate on the highest-income persons increases labour supply and enables the government to transfer more to low-income individuals.

Apps et al. (2014) demonstrate numerically that under a utilitarian social welfare function, the optimal piecewise linear income tax system can result in either a convex (increasing marginal tax rates) or a non-convex (decreasing marginal tax rates) tax structure. The latter will have a smaller guaranteed income under the optimal two-bracket piecewise linear tax system than under the optimal linear tax system, and the converse will hold in the former case. Which outcome occurs depends on the wage distribution and the elasticity of labour supply. When the wage distribution is Pareto and the elasticity of labour supply is the same for all types, marginal tax rates are higher in the top bracket than in the lower one. However, when the elasticity of labour supply is lower in the bottom eight wage deciles than in the top two, the marginal tax rate is higher in the lower tax bracket, which is intuitive. In contrast, when the wage distribution increases steeply at low wage rates and then follows the Pareto distribution at higher wages, the first-bracket marginal tax rate is now higher than the second bracket rate, and the transfer at zero income is higher than under a linear progressive tax. This case, which is the more realistic, has higher inequality at the bottom and lower at the top compared with the Pareto distribution.

Interdependent Utility Functions

The optimal income tax literature has been extended in a wide variety of ways, many of which we consider in the following two sections. In this section, we consider one extension to the social welfare function approach that raises both conceptual and practical issues. It involves the consequences for redistributive policy of utility functions being interdependent. Utility functions are interdependent when outcomes of one person directly affect the utility of another (rather than through prices). This can include altruism, envy, avarice, paternalism, or simply having preferences over another person's choices.

Choices resulting from interdependent utility functions may be observable, such as bequests, charitable donations, dependent care, or voluntary work. Interdependent utility can also influence the purchase of positional or status goods and the supply of labour to keep up with one's neighbours. Decision-making within families is particularly driven by interdependent utility. However, some forms of interdependent utility are not revealed by individual behaviour, including admiration or envy of one's neighbours or a distaste of their choices. Preferences over social outcomes are also not directly observable except through polling or voting choices.

A key issue is whether interdependent preferences should count for social welfare purposes. Some argue that all decisions reflect the revealed preferences of individuals, so should count as individual utility for social welfare purposes (Kaplow, 2008). Others argue that decisions made to fulfill interdependent utility preferences should not count in the decision-maker's social welfare function because that gives rise to double-counting (Hammond, 1987; Milgrom, 1993). For example, voluntary transfers would count as utility both to donors and to recipients. Our view, which we discuss in Boadway and Cuff (2015), is that counting the benefits of both donors and recipients in social welfare is questionable. Although the revealed-preference argument for double-counting is on the surface persuasive, there are compelling arguments against it. The benefit to a donor from the utility gains of a recipient should then apply to any form of interdependent utility whether revealed through transfers or not. Family members value the well-being of each other member, but there is no suggestion that this multiplicity of utilities be counted in social welfare. The same applies to feelings of altruism or avarice toward fellow citizens. Mirrlees (2007) notes the analogy between saving for one's own future and saving for heirs. He argues that we would not consider counting the utility we obtain now when saving for our future self. Some might regard the role of government redistribution as reflecting the altruism of the rich for the poor, and internalizing the free-riding from private donations (Hochman & Rodgers, 1969). There is no suggestion that in this case the rich taxpayers' altruism should be counted as social welfare in addition to the benefits of the transfer recipients. Voluntary transfers may not give utility to donors at all. They may represent voluntary transfers done out of a sense of obligation, making the donors worse off.

Decisions taken for interdependent utility motives have policy consequences. For example, voluntary transfers such as bequests or charitable donations affect both equity and efficiency. From an equity perspective, such transfers represent income to recipients that can be used for consumption and other purposes. For the donors, if the transfers are considered to represent utility benefits, they should be treated as consumption. However, if donors are regarded as forgoing utility when they make transfers, these should be deducted from their income. Efficiency consequences arise if one regards the transfers as yielding utility to the donors. If so, voluntary transfers simultaneously benefit the donors and the recipients, but the donors take into account only their own benefit. This constitutes an efficiency case for subsidizing voluntary transfers (Kaplow, 2008; Diamond, 2006). In practice, charitable donations are typically subsidized, while bequests are not, which is inconsistent. Potentially, the social welfare consequences of intra-family transfers are large. A transfer from one family member to another provides altruistic benefits to all family members so is like a public good. In principle, all such benefits of family transfers would have to be taken into account, which is potentially contentious, and in any case infeasible since most such transfers cannot be observed by the government.

Similar conceptual issues arise with the consumption of status or positional goods and with "rat race" labour supply. Most studies assume that the personal benefit from consuming status goods or from keeping up with the income of one's neighbour should not count as social welfare. Layard (1980) notes that if utility depends on status and income reflects status, income taxation is justified on both equity and efficiency grounds. It reduces excessive work due to status-seeking and increases social welfare. Frank (2005) discusses the negative consequences of utility depending on relative consumption. Consumers spend too much on "positional goods" (goods that signal consumers' position in a status ranking), which, rather than increasing their material well-being, simply competes to change their status ranking. Increasing taxation of these goods can be difficult, but increasing income tax progressivity reduces spending on luxury positional goods relative to more material goods, while leaving individuals in the same position as before. On the other hand, were status goods to count in social welfare, the tax system should be less progressive.

Some authors study the implications of interdependent utility for optimal income taxation. Oswald (1983) analyzes the Mirrlees optimal tax problem when individual utility functions include a term for either altruism or jealousy, and these terms count as social welfare. Marginal tax rates increase with jealousy and decline with altruism and can be negative in the latter case. Nishimura (2003) assumes that in addition to maximizing social welfare, the planner wants to restrict envy in the sense of individuals preferring others' consumption bundles. He adds limited envy as a constraint, which is binding for optimal income tax allocations in which high-income individuals are better off. This results in higher marginal income tax rates as well as higher commodity tax rates on goods preferred by low-wage individuals, which reduces the envy of the high-wage types.

Kanbur and Tuomala (2013) consider the case where individuals' well-being depends on their income relative to the average. Using theory and simulations, they show that as individuals put more weight on relative income, progressivity of optimal marginal income taxation also increases; that is, marginal tax rates rise faster with income. They also find that this increase due to relativity is dampened as pre-tax inequality increases. In any case, they find large negative average taxes at the bottom consistent with a sizable basic income. This average tax rate at the bottom becomes more negative as relativity increases, while marginal tax rates at the bottom increase as relativity increases.

The issue of interdependent utility functions has relevance for the form of basic income. If transfers to the poor are motivated by the altruism of the rich, the rich may well have preferences over the way the poor spend their transfers. This may lead the government to constrain the manner in which payments are made to the poor, such as by restricting their use to purchase food and other necessities rather than goods and services less valued by the rich. This would be inconsistent with transfers to the poor motivated by social welfare maximization, where the arguments of the social welfare function are utility levels as judged by individuals themselves.

Different Labour Market Environments

This section explores the relationship between optimal income taxation and basic income in more realistic labour market environments. In the standard optimal income tax literature, the demand for labour is perfectly elastic, so all labour supplied is employed at fixed wage rates. Once labour demand considerations are taken into account, wages may no longer be fixed, and employment may no longer be guaranteed. Changes in income taxes and transfers can affect market wages as well as employment, and this can change the optimal income tax structure, including the basic income guarantee. Both labour supply and labour demand responses become relevant.

It is useful at the outset to distinguish among various types of unemployment. We have already encountered voluntary unemployment in the form of either non-participation in the labour market or the choice of limited hours of work. The strength of these extensive and intensive labour supply responses to wage changes directly affect the optimal income tax structure through the efficiency effect. The voluntary unemployed may engage in non-market work, including care of children or the elderly and voluntary or charitable work. Transfers to the voluntary unemployed could be conditioned on their non-work activities. This has been neglected in the optimal tax-transfer literature, at least partly because of observability issues. Individuals may also choose not to participate in the labour market in order to pursue education or training, which are typically favoured by the income tax system.

Unemployment may not be voluntary. Individuals may be unable to work because of disability, maternity, or old age. Alternatively, they may be involuntarily unemployed—that is, willing to work but unable to find a job. Involuntary unemployment has an important temporal dimension. It may be temporary or frictional, such as when workers are temporarily laid off or are laid off but able to find work elsewhere. Unemployment may be more long term if workers are laid off because their skills are no longer in demand (structural unemployment), if they are unable to acquire a first job because of search frictions, or if there is excess supply of a given type of workers because of wage rigidities induced by minimum wages or efficiency wages. Individuals may be unable to find a job because they lack necessary skills.

A key issue is the relative size of transfers to each type of unemployed, and how they are delivered. The temporarily involuntary unemployed may be eligible for unemployment insurance, which can be based somewhat on insurance principles, so benefits are related to wages when employed. Maternity is also obviously temporary and triggers eligibility for unemployment insurance, while the elderly may be covered by public pension schemes. Disabled persons typically obtain targeted stand-alone transfers subject to some eligibility criteria. Transfers to the remaining categories would be based on income redistribution principles, and how much each group is entitled to depends on a value judgment. For example, how much should voluntary non-participants receive relative to the long-term unemployed who are unable to find work? This depends on how one treats persons with different preferences for work, as discussed in Fleurbaey and Maniquet (2011) and Ghatak and Maniquet (2019). Similarly, how should one treat those unable to work versus those who choose not to work, a distinction made in Cuff (2000).

An important consideration in determining transfers to different categories of persons is how the government obtains information about individuals. Transfers delivered through the income tax system are based on incomes reported to the tax administration voluntarily. To mitigate the possibilities of misreporting, audits are used, which may be random or triggered by some information reported. Those found to have misreported their incomes face a penalty. Transfers delivered outside the income tax system are generally administered by application where applicants are screened ex ante for eligibility (e.g., employment status, assets, disability), and monitored ex post for continuing eligibility. For example, welfare and unemployment insurance may be conditional on searching for jobs and accepting suitable job offers.

The monitoring of transfer recipients raises two concerns. First, there may be errors in monitoring, so some eligible applicants are found incorrectly to be ineligible, or vice versa. The possibility of inadvertent overpayment of transfers raises the difficulty of how, if at all, to recoup the overpayment. The possibility of errors can affect the optimal size of the transfers. Second, the need for monitoring can be intrusive and demeaning and can lead to stigmatization of potential applicants. Stigmatization can reduce the welfare of applicants and can deter some from applying for transfers. Welfare recipients can also find it difficult to find employment, especially in non-menial jobs. One of the arguments for a basic income that is based primarily on income and so can be administered through the income tax system is to avoid stigmatization. At the same time, ex ante monitoring may be better able to respond to unexpected changes in incomes than can self-reporting under the income tax system.[16]

We turn now to a more detailed consideration of how labour demand affects the optimal income tax structure, including the basic income guarantee. First, we discuss the case of endogenous wages. Then we turn to involuntary unemployment and distinguish between the cases where individuals are unable to work and where they are unable to find work, either permanently or temporarily.

Endogenous Wages

In the standard optimal income tax model, production is linear, so wage rates are fixed and determined by individual skills or productivity. In fact, wages can be endogenous for various reasons. They may be determined by standard demand and supply considerations. Stiglitz (1982) considered the case where relative wages of different skill types are determined in general equilibrium by their relative supplies. He analyzed the case of two skill types and argued that the marginal tax rate on the high-skilled type should be negative, rather than zero as in the standard analysis.

[16] The consequences of stigmatization for transfers is discussed in Jacquet and Van der Linden (2006). Different approaches to monitoring in transfer schemes may be found in Boadway and Cuff (1999, 2014), Kleven and Kopczuk (2011), and Boone and Lans Bovenberg (2013). Issues in designing transfer programs to deal with income volatility are discussed in Boadway et al. (2008).

Decreasing the top marginal tax rate encourages the supply of high-skilled workers, which in turn increases the relative wage of low-skilled workers, thereby enhancing redistribution. (Note that a low top marginal tax rate is consistent with a high average tax rate, which also encourages labour supply through an income effect.) While Stiglitz's focus was on the top marginal tax rate, similar considerations apply to low-skilled workers. A higher marginal tax rate at the bottom discourages low-skilled labour supply, which increases the low-skilled wage and enhances redistribution. Similarly, a higher participation tax on low-skilled workers in an extensive-margin model should increase the wage payment to low-skilled workers. Such reasoning suggests that increasing the basic income to the lowest-income workers will discourage their labour supply by an income effect and increase their equilibrium wage rate.

This effect of basic income on the wage rate of low-income workers is reinforced if basic income is taxed back. The latter entails a substitution effect that further reduces labour supply and increases the wage rate. Kasy (2018) argues that this favours a basic income guarantee with a positive taxback rate relative to the U.S. EITC system—or the CWB in Canada—where the effective marginal tax rate at the bottom can be very low, and possibly negative. Some of the benefit of the EITC accrues to firms whose wage payments are reduced. Rothstein (2010) simulates the effects of the EITC in the U.S. compared with a standard NIT system. He finds that a large proportion of the EITC accrues to firms through reduced wage payments. A $1 increase in EITC payments increases recipients' incomes by only $0.73. Moreover, workers not eligible for the EITC suffer a reduction in incomes due to the EITC-induced fall in wage rates. On the other hand, $1 spent on the NIT leads to an increase in recipients' incomes of $1.39 due to the induced increase in wage rates.

Another source of endogenous wages is work experience. Krause (2009) studies a simple two-period extension of the two-type optimal income tax model where second-period productivity and therefore wages are increasing in the number of hours worked in the first period. The government can impose a separate income tax structure in the first and second periods, an admittedly strong assumption. The optimal marginal tax rate on the high-wage workers is zero in both periods as in the standard model. However, a lower marginal tax rate—possibly negative—is imposed on low-wage workers in the first period. This encourages labour supply and increases their second-period wage rate, thereby both making low-skilled workers better off and increasing the ability to redistribute in the second period by relaxing the incentive constraint. Although Krause does not consider it, presumably a lower basic income guarantee in the first period will also encourage labour supply and learning-by-doing and therefore the second-period low-skilled wage rate.

Wage rates may also be affected by human capital investment, including education. Lans Bovenberg and Jacobs (2005) and Jacobs and Lans Bovenberg (2010, 2011) have incorporated human capital investment into optimal income tax models. They show that a progressive earnings tax discourages human capital investment since the marginal tax rate is higher on increases in earnings than on forgone earnings resulting from such investment. This effect can be mitigated by subsidizing human capital investment or taxing capital income. Stantcheva (2015) studies the case for subsidizing human capital accumulation in a dynamic optimal income tax analysis. The analysis is complex and shows that the argument for subsidizing individuals' human capital investment depends on the extent to which a subsidy increases wages and therefore labour supply, whether investment in human capital is complementary with labour supply (learning-and-doing) or substitutable (learning-or-doing), and the differential effect of training on persons of different skill levels. Blumkin and Sadka (2005) analyze the effects of intergenerational mobility on tax progressivity when individuals' wage rates depend on their parent's wage rate and on the level of education, and the government levies a linear progressive tax. They show that optimal tax progressivity falls with mobility, and with inequality if mobility is high enough. These results have limited implications for basic income except to the extent that basic income encourages education or training.

Involuntary Unemployment I: Those Unable to Work

Two strands of the literature have addressed the optimal transfers to those unable to work. One is the tagging literature that assumes the population can be divided between those able to work and those unable to work or disabled and explores the integrated tax-transfer treatment of both, given an assumption about relative social welfare weights to the two groups. The other emphasizes differences in preferences rather than disability as a determinant of voluntary unemployment and considers their implications for optimal redistribution policy.

Tagging

Some individuals may be unable to work or may be able to work only a limited amount due to disability or social conventions. The optimal transfer to the disabled relative to those who are able but choose not to work depends upon a value judgment. In practice, the disabled are typically eligible for a higher transfer than the voluntary unemployed, assuming they can be identified. For one thing, disabled persons may require a higher income to attain the same level of utility as those voluntarily unemployed given that they face a higher cost of living.[17] As well, transfers to the voluntary employed may be held below the first-best optimal level in order to discourage able workers from choosing not to work. If the government can observe who is disabled, the incentive constraint that precludes the able from choosing not to work will not apply, and the disabled can be given a higher transfer. If the government cannot identify who is disabled, it must make the same transfer to those unable to work and the voluntary unemployed. Recall that this is the assumption in the extensive-margin optimal income tax approach discussed above.

One way to address this information issue is to assume that the government can acquire some information that enables the incentive constraint to be relaxed. One approach, following Akerlof (1978) and Parsons (1996), assumes that disability is imperfectly observable. An observable signal or *tag* exists that is positively correlated with disability. Transfer recipients can be divided into tagged and untagged groups, with those tagged having a higher proportion of disabled relative to able persons. Within each group, transfers to those not working are limited by an incentive constraint, but lump-sum transfers can be made between groups. Under a utilitarian social objective, transfers are made from the untagged to the tagged group until the average marginal utility of income in each group is equal. However, intergroup redistribution creates inequality: the untagged disabled are made worse off (type-I statistical errors), while the tagged able become better off (type-II errors). If aversion to inequality increases, the transfer from the untagged to the tagged group falls, reducing the value of tagging. In the limit, with a maximin objective, tagging would be of no use since it makes some of the disadvantaged worse off.

The value of tagging is limited by other factors besides aversion to inequality. The less accurate is tagging—especially the more type-I errors—the less valuable it is. Some may object to tagging on horizontal equity grounds since otherwise identical persons are treated differently depending on whether they are tagged, though for others social welfare overrules horizontal equity concerns (Kaplow, 2001). The value of tagging may be further compromised if a stigma attaches to being tagged (Jacquet & Van der Linden, 2006). The stigma may reflect the shame of being identified as tagged, even if only the tagging administrator observes it. The stigma may also result from type-II errors. Public knowledge that some non-deserving persons receive transfers may throw suspicion on all transfer recipients. The failure of needy persons to be tagged may reflect low

[17] However, Sen (1973) has pointed out that a utilitarian government may choose to transfer less to the disabled than to the able unemployed. It will seek to equalize the marginal utility of consumption among transfer recipients, and that may entail higher transfers to those who are most able to convert income into utility.

application or take-up rates due to cost or complexity. Complexity may even be used to discourage non-deserving applicants from applying (Kleven & Kopczuk, 2011). Finally, agency problems between the government and tagging administrators may compromise the accuracy of tagging. Social workers may exert low effort if it cannot be detected (Boadway et al., 1999). They may be more or less sympathetic to applicants than the government, and there may be no natural mechanism to ensure that social workers who share the government's values choose to become social workers (Prendergast, 2007).

The tagging literature was inspired by the desire to enhance the ability of government to overcome informational constraints on targeting transfers to deserving persons. The principle has been taken up in the optimal income taxation. As tagging suggests, some observable characteristic may allow policy-makers to divide taxpayers into identifiable groups, each with different skill distributions. Characteristics include age (Banks & Diamond, 2010; Weinzierl, 2011), gender (Cremer et al., 2010; Alesina et al., 2011), height (Mankiw & Weinzierl, 2010), parental status (Bastani et al., 2019), and state/province of residence (Gordon & Cullen, 2012). The population is divided into two or more identifiable groups, and an optimal income tax applies within each group, with intergroup redistribution. As long as the distribution of skills differs across groups, social welfare can be enhanced by tagging, possibly even if the social welfare function is maximin (Boadway & Pestieau, 2007). The properties of income taxes in each group vary with the specific application. Immonen et al. (1998), using simulation techniques, find that marginal tax rates increase with skills in the group with higher average skills, but decline in the other group. Boadway and Pestieau (2007), using a two-type example, show that the income tax should be more progressive in the group with higher average skills. Cremer et al. (2010), who tag by gender, find in a model calibrated to U.S. data that higher-skilled males would lose from tagging. Some may object to tagging by gender or other characteristics on the grounds of discrimination. Banks and Diamond (2010) argue that this objection does not apply to age, since everyone experiences different ages over their lifetime.

Different Preferences

In the standard tagging literature just discussed, individuals with different characteristics share the same preferences. More challenging issues arise when they differ by preferences, especially preferences for leisure. Two difficult problems arise. First, preferences are not directly observable, though they affect behaviour. The population cannot be divided into two groups to which intergroup lump-sum transfers can apply. At best, the government may elicit such information at a cost by monitoring (Boadway & Cuff, 1999) or by separating contracts (Blackorby & Donaldson, 1988). In the optimal income tax literature, the government cannot identify persons with different preferences but can condition the tax system on knowledge it has about the distribution of preferences.

Second, aggregating individuals with different preferences into a social welfare function is conceptually problematic. The pattern of optimal marginal tax rates can vary significantly according to the weight put on different preference groups (Cuff, 2000; Boadway et al., 2000; Choné & Laroque, 2010). If individuals with a low preference for leisure are given a high weight, the optimal tax system may redistribute from lower- to higher-income persons at the bottom end, resulting in negative marginal tax rates. A low basic income will reflect this, if a high proportion of those choosing not to work have a high preference for leisure. The problem is compounded by the fact that some persons earning little or no income are unable to work due to disability, but might nonetheless have a strong preference for working if they could.

An alternative way of dealing with heterogeneous preferences, as mentioned earlier, is not to reward or punish individuals based on their preferences. The value judgment invoked is that policy should aim to redistribute among individuals to compensate only for differences in characteristics over which they have no control, such as wage rates or skills—the *principle of compensation*. At the same time, they should neither be rewarded nor penalized for characteristics over which they have control—the *principle of responsibility*. Individuals are typically assumed to have control over their preferences (Fleurbaey & Maniquet, 2011), though with some limits resulting from their socio-economic backgrounds (Roemer, 1998). Despite the attractiveness of these two principles, they cannot generally both be satisfied at the same time. For example, suppose there are two wage types (high- and low-wage) and two preference types (high and low taste for leisure), and assume the government can observe each individual's wage–preference type. For each preference type, utilities should be equalized by lump-sum transfers (the principle of compensation); for each wage type, individuals of different preferences should get the same transfer (the principle of responsibility). It is straightforward to show geometrically that these two policies cannot be satisfied at the same time (Boadway, 2012, p. 214).

Fleurbaey and Maniquet (2011) take the approach that while the two principles cannot both be satisfied, one could be fully satisfied and the other partially. Consider the case where the principle of compensation is fully satisfied, and the principle of responsibility only partially. Full compensation is achieved by an equal-preferences transfer, which is a variant of the Pigou–Dalton transfer axiom in income distribution theory. Imagine two individuals with the same preferences and different wage rates supplying the same labour but obtaining different consumption. By the Pigou–Dalton axiom, a transfer that reduces their consumption inequality while holding labour supply constant improves social welfare. Fleurbaey and Maniquet capture partial responsibility by laissez-faire selection, which says that if all individuals had the same skills but different preferences, laissez-faire would be socially optimal. The social welfare ordering must satisfy laissez-faire selection and the Pigou–Dalton transfer principle, along with the Pareto principle and the independence of irrelevant alternatives. Fleurbaey and Maniquet show that this implies a maximin social ordering with the least well-off individuals being those with the lowest wage rate and the weakest preferences for leisure. The optimal income tax system maximizes their well-being subject to satisfying incentive constraints. In this case, the least well-off persons will not have the lowest incomes. As Ghatak and Maniquet (2019) note, a low basic income will apply to low-wage persons with a high preference for leisure.

Different characterizations of partial responsibility yield different results. An example that Fleurbaey and Maniquet (2011) consider is the requirement that transfers that reduce differences in the budget line of two persons with identical wage rates lead to increases in social welfare if the person with the higher budget constraint has the least preference for leisure. Under this assumption, the individual with the lowest wage and the highest preference for leisure is considered the worst off. In this case, the basic income guarantee accruing to the persons with the lowest income would be correspondingly higher. Fleurbaey and Maniquet also consider cases where responsibility is full and compensation partial. Their approach emphasizes the need to establish a set of principles that determine who is the least well-off and therefore most deserving of transfers. As a methodology for addressing differences in preferences, it is useful, but it provides limited guidance from a policy perspective. While the least well-off persons include those with the lowest wage rate, whether the most deserving are those with the lowest or highest preference for leisure depends on the assumptions adopted. Moreover, no attention is given to the optimal size of the transfer to those unable to work relative to those with a high preference for leisure.

Involuntary Unemployment II: Those Not Finding Work

So far, the labour market models described assume that the demand for labour satisfies the supply of labour at the market wage. In the case of intensive-margin models of employment, individuals choose the hours they work based on their preferences over consumption and leisure, while in the extensive-margin case individuals choose to work if the after-tax wage they would receive is high enough. Apart from those unable to work, all unemployment is voluntary.

In this subsection, we explore labour market environments where individuals may be willing and able to work or to work more hours at the prevailing market wage, but opportunities are not available to them. In these environments, some workers are either involuntarily unemployed or underemployed. Much of the literature focuses on the former, reflecting the assumption in extensive models of labour supply that workers cannot change their hours of work. Involuntary unemployment can result from labour market inefficiencies, such as matching frictions in job search or efficiency wages. It can also result from minimum wages that cause excess labour supply.

We consider the implications of involuntary unemployment for the optimal tax-transfer system, including the basic income guarantee. There are two overriding effects. First, when all unemployment is voluntary, the optimal transfer to unemployed individuals is constrained by the behavioural response it elicits, such as a reduction in participation or hours of work. For individuals who are willing to work but for some reason unable, this behavioural response is mitigated, and an income transfer will not change their labour earnings. The more of those not working who are involuntarily unemployed, the less will be the labour market consequences of increased transfers to the unemployed. Changes in the optimal income tax system will affect both the extent of unemployment and the wage obtained by those who have a job. Second, with involuntary unemployment, a basic income takes on an insurance role in addition to a redistributive one. Under these conditions, the optimal basic income is larger than in the case with only voluntary unemployment. The longer the spell of involuntary unemployment, the more the basic income plays a redistributive rather than an insurance role, implying it is important to make a distinction between long-term and short-term unemployment.

The policy response to short-term unemployment is unemployment insurance, whereas policy to address long-term unemployment emphasizes redistribution. Temporary layoffs or frictional unemployment due to job turnover are covered by unemployment insurance and self-insurance. Structural unemployment leading to persistent earnings loss is addressed in part by redistributive policies. Retraining typically results in only partial recovery of earnings loss (Jacobson et al., 1993, 2005). Displaced workers who suffer permanent wage rate declines are comparable to those who have low skills at the outset, so are candidates for optimal income tax policies. An alternative approach to address wage shocks is wage insurance (LaLonde, 2007), though there is limited literature on that.

The normative policy literature has tended to draw a sharp distinction between workers who face permanent unemployment and ought to receive redistributive transfers versus those who are temporarily unemployed and ought to receive unemployment insurance. This is so despite the fact that the same search or matching models may be used to generate unemployment in both cases. We follow the literature and deal with permanent and temporary unemployment separately in the following subsections.

Long-Term Involuntary Unemployment

In the case of the long-term involuntary unemployed, a basic income functions mostly as a redistributive tool for individuals who are unable to find a job due to frictions in the labour market that result in excess labour supply. Optimal income tax models typically adopt classic dynamic search or matching models (Pissarides, 2000; Rogerson et al., 2005) to explain involuntary unemployment. For simplicity, static counterparts of these models are typically used, in which individuals make one labour market choice that has permanent consequences. There is limited use of other models of involuntary unemployment, such as efficiency wages or turnover models, in the optimal income tax context.

The key elements of the static search or matching models used are as follows. Firms seeking to hire workers of a given skill type post vacancies, where there is a fixed cost per vacancy posted. Once hired, workers produce output according to a concave production function. Firms are competitive and are free to enter or exit industry. A zero-expected profit condition determines the total number of vacancies posted V. Unemployed workers of each skill type choose to look for jobs for their type. The total number of workers of a given type searching is U. A matching function, $M(U,V)$, with $M_U(U,V) > 0$ and $M_V(U,V) > 0$, determines the number of positions filled. The matching function is often assumed to be linear homogeneous, so $M(U,V) = Vm(1/\theta,1)$, where $\theta = V/U$ is an index of market tightness. Jobs are filled randomly, so the probability of filling a job is given by $M(U,V)/V = m(1/\theta,1) \equiv \pi(\theta)$, where $\pi(\theta)$ is decreasing in θ. The probability of a worker obtaining a job is $M(V,U)/U = \theta\pi(\theta)$, which is increasing in θ.

The wage payment y—or earnings, since labour supply is fixed—is determined by bargaining between the worker and the firm after a match has been obtained. (Alternatively, competitive wage-setting will yield similar results, where firms post wages that they will pay for filled vacancies.) With Nash bargaining, the equilibrium wage maximizes the Nash product of the surpluses of workers and firms, $\big(y - T(y) - b\big)^{\rho}(a-y)^{1-\rho}$, where $T(y)$ is the earnings tax function, b is the transfer to the unemployed, a is the productivity of the worker, and ρ is the relative bargaining power of workers.[18] When workers differ in skills, it is often assumed that this wage bargaining process applies separately to each skill type. While this simplifies the analysis of the optimal tax-transfer system, it restricts workers to searching for jobs of their skill type.

Hungerbühler et al. (2006) were the first to analyze the consequences for optimal income taxation of involuntary unemployment due to search frictions. They adopt an extensive-margin participation model, in which workers choose whether to seek a job at their skill level. Those who do so engage in a job search, and their success depends on a skill-specific matching technology of the above sort. The value of leisure is the same for all regardless of skills, and this leads to a cut-off skill level $\tilde{a}$ such that workers participate if and only if $a \geqslant \tilde{a}$. For those who obtain a job, wage earnings y are determined by Nash wage bargaining, which is efficient since the Hosios condition is satisfied. The government observes earnings y and chooses a nonlinear income tax $T(y)$ to maximize a social welfare function with inequality aversion subject to a budget constraint and incentive constraints. Since the government cannot distinguish the voluntary from the involuntary unemployed, it pays the same transfer b to all. While the government observes employment earnings y, it cannot observe worker abilities a. This implies that firms employing workers of one skill level can mimic the bargaining outcomes of others. To preclude this, an incentive constraint applies to wage bargains, which along with the endogenous participation decision constrains government redistribution policies.

The tax structure $T(y)$ affects both bargaining outcomes y and the cut-off skill level $\tilde{a}$ determining participation. In particular, participation is decreasing in the average tax rate $T(y)/y$. At the same time, the tax affects the wage bargain in conflicting ways. An increase in total tax liabilities $T(y)$ reduces workers' surplus, thereby causing an increase in wages and a decrease in employment. An increase in the marginal tax rate $T'(y)$ reduces the incentive for workers to bargain for higher wages, so wages fall and employment rises. The participation tax applying to the marginal participant is always positive, unlike in Saez (2002). The marginal income tax rate $T'(y)$ is positive for all y, while the average tax rate $T(y)/y$ is increasing. Numerical simulations indicate that marginal tax rates are much higher than in the Mirrlees intensive-margin model.

[18] Suppose the matching function takes the Cobb-Douglas form, $M(U,V) = U^{\alpha}V^{1-\alpha}$. Hosios (1990) showed that if $\rho = \alpha$, bargaining is efficient in the sense that the externalities of search are internalized. The Hosios condition implies that the share of the surplus captured by workers in the bargaining process reflects workers' relative productivity at generating matches, and similarly for firms, so that workers' search effort and firms' decisions to create vacancies are efficient. Then the wage under Nash bargaining maximizes the expected surplus of workers, $\theta\pi(\theta)(a - T(y) - b)$.

Hungerbühler et al. (2006) make some fairly strong assumptions. They assume a uniform value of leisure for all workers, implying that all workers with skill above the cut-off participate, while none of those below do. Bargaining to determine earnings is skill-specific and efficient, and firms and their workers are able to mimic wage bargains of other skill levels. Workers decide only whether to participate, and search unemployment is permanent. The government cannot distinguish the voluntary from the involuntary unemployed, which seriously restricts its ability to make transfers. Some of these assumptions have been relaxed in subsequent literature.

Lehmann et al. (2011) allow heterogeneity in the value of leisure at all skill levels. This leads to both voluntary and involuntary unemployment for all skill types. Variable participation at each skill level creates a participation elasticity as in Diamond (1980) and Saez (2002), which both is realistic and makes comparison with the latter possible. Lehmann et al. first study the maximin case, where analytical results are possible. The least well-off are the involuntary unemployed, who all receive the same transfer. Optimal marginal tax rates are positive everywhere and tend to be higher than in models with no involuntary unemployment. Participation tax rates are positive everywhere, as in the pure extensive-margin case under maximin. If the elasticity of participation falls with skill level, which they suggest is plausible, the average tax rate is increasing in earnings, so is progressive. Simulations are needed when social welfare exhibits finite inequality aversion. In this case, the participation tax and the marginal tax rate can be negative at the bottom, as in Saez (2002).

Jacquet et al. (2014) analyze a variant to Lehmann et al. (2011) that differs in two main ways. First, they preclude the possibility of one skill level mimicking another by choosing the same wage. The government can observe wage bargaining in each group, so the incentive constraint no longer applies. Second, wage bargaining takes the form of Kalai bargaining rather than Nash bargaining. This implies that the shares of the surplus accruing to workers and firms are exogenously fixed. Earnings of a type-a worker are $y_a = \rho_a a + (1-\rho_a)\big(T(y_a) - b\big)$, where ρ_a is the proportion of the surplus from the bargain accruing to workers,[19] so the equilibrium wage is increasing in the level of the tax and does not depend on the marginal tax rate. An increase in the tax on a given skill level both reduces labour demand—since y increases—and reduces labour supply—since participation decreases.

Using this, the optimal participation tax for workers of skill level a is a straightforward

[19]More precisely, the wage rate y for type a maximizes the minimum of $(y - T(y) - b)/\rho_a$ and $(a-y)/(1-\rho_a)$, so ρ_a is the bargaining strength of the worker.

extension of (21) for the pure extensive-margin case:

$$\frac{T(y_a)+b}{c_a-c_0}=\frac{1-g_a\rho_a\left(1+\eta_a^D\right)}{\gamma\left(\eta_a^D+\eta_a^P+\eta_a^D\eta_a^P\right)}$$

where η_a^P is the elasticity of participation of a type-a worker, and η_a^D is the elasticity of demand for a type-a worker with respect to surplus $a-y_a$. This makes it clear that there are both demand and supply influences at work in labour-matching models. Jacquet et al. (2014) simulate the optimal values of participation tax rates using values for the elasticities η_a^P and η_a^D that are empirically reasonable. They find that optimal taxes are systematically lower than in pure extensive-margin models without involuntary unemployment, which is intuitive since there is now an additional decision margin influenced by the tax. They also find that the optimal employment tax decreases when the elasticity of labour demand increases relative to the participation elasticity.

Kroft et al. (2020) build on Lehmann et al. (2011) by developing a sufficient statistic approach to formulating the tax structure. Individuals face a distribution of occupation-specific discrete search costs and decide whether to search for work in a particular occupation. As in Saez (2002), wages/output are observable, and income tax liabilities are occupation-specific. Search activity is unobservable, so the transfer to both the voluntary unemployed (those not searching for work) and the involuntary unemployed (those unable to find work) are the same. How wages and employment are determined is not specified. Rather, it is simply assumed that the equilibrium distribution of wages and employment across occupations depends on the government's tax policies. Employment in an occupation equals the product of the number who search for a job in that occupation and the conditional probability that a participant will find a job.

Kroft et al. (2020) differentiate between micro responses to taxes, captured by changes in participation or, equivalently, search decisions, given wages and conditional probabilities of employment, and macro responses to taxes, captured by the general equilibrium changes to wages and conditional probabilities of employment. Participation decisions depend on expected utilities that in turn depend on taxes through both micro responses and macro or general equilibrium responses (through changes in wages and conditional probabilities). The optimal tax formula depends on both micro and macro responses and is expressed using sufficient statistics.

This can most easily be seen in the pure extensive-margin case where individuals search for work in only one occupation and the macro effects of a tax change are restricted to effects of the wage and conditional probability within a given occupation. The optimal participation tax rule is a straightforward extension of (21) above:

$$\frac{T_i + c_0}{c_i - c_0} = \frac{1 - g_i \frac{\pi_i}{\pi_i^m}}{\eta_i}$$

where the new terms on the right-hand side are the macro participation elasticity π_i, the macro employment elasticity η_i, which captures changes in both participation and conditional probabilities, and the micro participation elasticity π_i^m.

Kroft et al. (2020) define g_i as the social welfare weight of giving individuals in occupation i one more dollar, holding wages and conditional probabilities fixed, as in Saez (2002). Since this additional dollar results in general equilibrium effects, g_i is multiplied by the ratio of the macro to micro participation elasticities. With fixed wages and full employment case, the participation tax reduces to (21), since with fixed conditional probabilities of employment η_i will be equal to π_i, and with fixed wages, the micro and macro participation elasticities will be the same.

Having a negative participation tax rate at the bottom of the earnings distribution now depends on whether g_1 is greater than the ratio of the micro to the macro participation elasticity (π_1^m/π_1). The size of this ratio depends on the assumed underlying wage determination process and could be greater or less than one. For example, in the case of risk-neutral individuals and the search/matching model discussed above, the ratio will be greater (less) than one if the workers' bargaining power is less (greater) than the elasticity of the matching function, and will be equal to one if the Hosios condition is satisfied. Notably, the equity effect of an increase in the tax rate depends on general equilibrium responses of wages and conditional probabilities, while for efficiency what matters is the macro response of employment.

Kroft et al. (2020) extend the model to allow for a form of "intensive" labour responses as in Saez (2002), in which individuals can choose to work in one of two adjacent occupations. In this case, changes in taxes in one occupation can have spillover effects on wages and employment in other occupations, and the optimal tax rule is modified accordingly. Simulations of the optimal income tax system in this mixed model illustrate the effect of these macro responses on the transfer to the unemployed and the marginal tax rates. First, assuming zero spillover effects and equal micro/macro participation elasticities, they show that using their lower estimated macro employment elasticities rather than the micro employment elasticity results in a higher transfer to the unemployed. In both cases, the employment tax for the lowest-earning occupation is negative, and marginal tax rates are negative at low incomes. Then, using the macro employment elasticity, they show that for macro/micro participation elasticity ratios greater (less) than one, the employment tax for the lowest-earning occupation is negative (positive) and the transfer to unemployed is smaller (larger). Further, when this ratio is less than one, marginal tax rates are positive for all income levels.

Importantly, their results suggest that the transfer to the unemployed will be greater when incorporating these general equilibrium responses, and this impact will be greater when the government has lower redistributive tastes. Despite the uncertainty surrounding the estimates and values of the sufficient statistics used, they conclude that endogenous wages and employment pushes the optimal tax system toward a negative income tax system and away from an EITC.

The papers discussed so far assume that the government does not distinguish between the voluntary and involuntary unemployed, and makes the same transfer to each. Boadway and Cuff (2014) assume that the government can observe imperfectly those who are involuntarily unemployed using an unspecified monitoring mechanism. The model is similar to Lehmann et al. (2011), with some exceptions. Individuals vary in both skill and cost of working. Those who choose to participate can either opt for a job of their skill level or enter the next-lower-skilled job market, where they incur a lower disutility of work than low-skilled workers. Higher-skilled individuals have a higher probability of finding a job. The government gives a different transfer to the voluntary and involuntary unemployed, and this gives rise to three incentive constraints: a participation constraint, a constraint discouraging high-skilled workers from choosing low-skilled jobs, and a constraint discouraging the voluntary unemployed from pretending to be involuntary unemployed. The government uses costly and imperfect monitoring to identify the involuntary unemployed. In the optimum, the high-skilled workers with the lowest disutility of work will search for high-skilled jobs, while those with the highest disutility of work will not participate, and those with intermediate disutility levels will search for low-skilled jobs. More intensive monitoring increases the difference between transfers to voluntary and involuntary unemployed and increases the transfer to the voluntary unemployed.

In the optimal income tax literature, search or matching frictions are the source of involuntary unemployment. In principle, involuntary unemployment could be a result of firms paying efficiency wages. In the efficiency wage model of Shapiro and Stiglitz (1984), firms offer workers wages above the market-clearing level to induce them to work harder—that is, not to shirk. Workers who get caught shirking are laid off and receive unemployment transfers. The incentive not to shirk is increasing in the difference between the wage paid and the transfer received if laid off. Offering higher transfers to the unemployed tightens the no-shirking constraint. Firms have to offer higher wages to discourage shirking, and more unemployment is generated. The result is a limit on transfers to the unemployed, or to the basic income.

Another source of involuntary unemployment is excess labour supply as a result of minimum wages. Much of the focus of the optimal tax literature has been on determining whether a minimum wage can be a welfare-improving policy alongside the optimal nonlinear income tax. Boadway and Cuff (2001) demonstrate in the standard nonlinear income tax model that if the government can perfectly monitor job search/offers and requires all job offers to be accepted, then with a binding minimum wage, only individuals with abilities greater than the minimum wage will be offered jobs. Consequently, the government has information about the abilities of those who are not working and can provide a higher transfer to them. Other papers have focused on economies where wages are determined endogenously. In such a setting with intensive-margin labour decisions, Marceau and Boadway (1994) demonstrate that a minimum wage inducing involuntary unemployed of low-skill workers can improve social welfare if the low-skilled have a negative tax liability. Lee and Saez (2012) obtain a similar result in a model in which individuals make participation decisions and differ in their cost of working in one of two possible occupations. They assume efficient rationing so that unemployment induced by a minimum wage will affect those with the highest costs of working. Then, a minimum wage in the low-skilled occupation can be welfare-improving if participation in this occupation is being subsidized. A minimum wage can mitigate positive labour supply/participation distortions induced by the optimal income tax system and enhance redistribution to the low-skilled.

Other papers have considered settings in which there is temporary involuntary unemployment, and a minimum wage policy can have redistributive and efficiency-enhancing roles. In Hungerbühler and Lehmann (2009), there is wage bargaining between workers and firms, and the bargaining share of workers is assumed to be less than the elasticity of the matching function, so the Hosios condition is violated, and bargaining is inefficient. A binding minimum wage at the bottom of the wage distribution can be welfare-improving. Similarly, Lavecchia (2020) shows that a minimum wage can reduce search congestion externalities. He demonstrates that without search frictions, if the optimal participation tax is negative for the lowest-skill occupation, then imposing a minimum wage offsets this distortion and results in greater redistribution. The upshot of this discussion is that the minimum wage can be a complementary policy tool for a basic income.

The models of this subsection assume that the absence of a job match gives rise to permanent unemployment, so the policy remedy involves redistribution. We now assume that matching-induced unemployment is temporary: workers who do not find matches in the current period expect to do so sometime in the future, although the duration will be uncertain. In this context, the government policy response is to provide unemployment insurance, assuming that such insurance cannot be provided privately or through self-insurance. Given that unemployment is temporary, completeness dictates that explaining job loss becomes part of the analysis. In much of the literature, either this is suppressed, or job loss is assumed to be random.

Temporary Involuntary Unemployment

When involuntary unemployment is temporary, its effects on human capital and future wages are limited, and a basic income functions more as a form of unemployment insurance than a redistributive transfer (although it is also that). A targeted basic income would act as an income-tested form of unemployment insurance, where workers with lower wealth and lower wages—and therefore less ability to self-insure—would receive a higher transfer. The same search models used in the permanent unemployment case are used in the temporary unemployment case, except they must be dynamic models, where individuals make choices each period and can transition from one job to another. In such models, agents take into account the possibility of future matches. Depending on the model, this may lead them to choose to hold off accepting a job if they believe a better one can be found by waiting. Alternatively, in models where it is relatively easy to find a job once you already have one, individuals may accept jobs they would not in a static search model. In dynamic search models, an unemployment transfer can affect things like job turnover, the duration of search, and search intensity.

The literature on optimal unemployment insurance has largely ignored any redistributive role and focused on efficiency concerns. Although studies of the design of unemployment insurance have been around a long time (as surveyed in Karni, 1999, and Topel and Welch, 1980), the recent literature has been inspired by Chetty (2008). He considers a representative-individual model, in which the individual faces some chance of becoming unemployed. The individual can respond by searching for work if unemployed and self-insuring against the income loss due to unemployment. In the absence of full self-insurance (or a private unemployment insurance market), government-provided unemployment insurance serves to smooth consumption between employment and unemployment but discourages search. Optimal unemployment insurance trades off these two effects. Chetty considers both a static model and a dynamic one in which duration of unemployment spells depends on search.

The efficiency analysis of unemployment insurance has been extended in a number of directions. Chetty and Looney (2006) argue that liquidity constraints are tighter for workers with low family income and wealth, and the latter face higher income volatility. Targeting unemployment insurance to them improves efficiency by avoiding moral hazard effects for higher-income workers. Crossley and Low (2010) find that liquidity constraints (the ability to self-insure) fall with age, so unemployment benefits should be lower for older persons. Cremer et al. (1996) study optimal unemployment insurance benefits when duration of unemployment is variable and workers face the possibility of a job mismatch. Benefits increase the duration of unemployment and allow some workers to wait for the right offer. Others, however, end up permanently unemployed, and benefits are regressive. Optimal benefits trade off these effects. Landais et al. (2018) augment the Chetty (2008) model by assuming unemployment is determined by a standard matching technology with wages determined by bargaining. With matching models, labour market tightness (the ratio of vacancies to job-seekers) is endogenous and may be inefficient because wage bargaining is inefficient. In this context, optimal unemployment insurance involves not just a trade-off between insurance and search effort, but also a correction for inefficient market tightness. If labour market tightness is too low because bargaining leads to excessive wages, unemployment insurance should be lowered to increase tightness, and vice versa. Landais et al. argue that since tightness falls during recessions, unemployment insurance should rise. Spinnewijn (2015) finds that the unemployed overestimate their chances of getting a job, so they search and save too little. To correct for these behavioural biases, unemployment benefits should be appropriately higher.

These models assess unemployment insurance using efficiency criteria. Boadway and Cuff (2018) consider an extensive-margin optimal income tax setting with search unemployment, in which unemployment insurance transfers exist alongside transfers to the voluntary unemployed. Their model is similar to that of Lehmann et al. (2011) with several exceptions. Workers are risk-averse but are unable to insure themselves against income losses if involuntarily unemployed. They can choose to be voluntarily unemployed or to search for a job, in which case they decide how much search effort to exert. The government observes whether workers search or not but cannot observe their type if unemployed. For those who obtain a job match, the government observes both the output they produce and the wages they obtain by bargaining. The government imposes a piecewise linear income tax such that workers of different wages face different marginal tax rates and lump-sum transfers. As well, the government sets a separate transfer to the voluntary and involuntary unemployed, where the latter is equivalent to an unemployment insurance benefit.

The Boadway–Cuff model includes both insurance and redistribution concerns, possible inefficiencies of search and wage-setting, and individual participation and search effort decisions. Government social welfare maximization leads to assigning different policy instruments to different policy targets. Marginal wage tax rates affect bargained wage rates and are chosen to achieve efficient wage-setting. This implies that optimal choices of transfers to employed workers of different types (i.e., average tax rates) and to the voluntary and involuntary unemployed need not take account of their effects on wage rates. The optimal unemployment benefit trades off insurance with participation and search incentives but is independent of equity weights. Optimal redistribution is achieved by type-specific transfers (average tax rates) to the employed and transfers to the voluntary unemployed and is constrained by both participation and search decisions. Overall, these results provide some justification for separating unemployment insurance design from redistributive tax-transfer policy.

The requirements of the analysis are, however, demanding, and more work needs to be done. The assumption that the government can observe who is searching for a job so can distinguish the voluntary from the involuntary unemployed allows it to target redistributive transfers to the former and unemployment insurance to the latter. As well, the absence of long-term unemployed absolves the government from differentiating transfers to the disabled and those unable to find work. In practice, governments impose eligibility requirements on job search that can only imperfectly be enforced by monitoring. The assumption that redistributive taxes can be perfectly targeted to workers with different wage rates or earnings and that participation is the only choice workers make ensures that redistribution can be assigned to the income tax system, leaving unemployment insurance to address insurance. In the real world, policy instruments cannot be so finely targeted. Nonetheless, the analysis makes the useful distinction between insurance and redistribution functions of taxes and transfers and suggests that in the context of a basic income guarantee, unemployment insurance and basic income might be viewed as complementary policy instruments.

Behavioural Economics Considerations

Behavioural economics studies economic outcomes and policies when individuals do not behave as rational self-interested agents (Bernheim et al., 2018, 2019). Individual decisions might exhibit present bias if they put excessive weight on current versus future benefits or costs. Present biases lead to undersaving, procrastination, and unhealthy consumption choices. Individuals may be hampered by bounded rationality and rely on rules of thumb, heuristics, or other arbitrary decision rules. They may be inattentive and not take proper account of available information, or their choices might be influenced by context irrelevant to the underlying economic decision. They may also make choices against their own interest, because they are motivated by moral considerations or are abiding by social norms.

How governments should design policy interventions in light of these behavioural considerations has received a great deal of recent attention (Congdon et al., 2011). Governments may impose taxes, subsidies, or regulations to influence irrational choices (e.g., tobacco and drug consumption, vaccinations, mandatory savings) or regulate the markets for products that are complicated for consumers to understand (e.g., financial instruments, pharmaceuticals). They may design policies to "nudge" consumers into making the best choices (Thaler & Sunstein, 2008). An important conceptual issue that arises in designing policies in light of behavioural considerations is the legitimacy of government overriding the preferences of individuals to correct their behaviour. Our concern is how behavioural economic considerations influence optimal income tax policies. There is a limited literature on this, and it typically regards government intervention as legitimate on the grounds that government is taking the long-run interests of individuals into account.

We begin with a brief discussion of social norms, since these influence individual behaviour in ways that are potentially important for redistribution policies, including basic income, and then consider the design of optimal linear and nonlinear income taxes in light of behavioural biases.

Social Norms

Social norms can be understood as the propensities of individuals to take actions out of an intrinsic desire to follow the social norms independent of individual utility, and the propensities of which are shared by a population (Elster, 1989). Individuals may also make decisions on whether or not to follow one social norm or another (or with regard to how closely they follow them) depending on how other individuals appear to be behaving. This has been shown to be particularly important in the context of tax compliance.[20] One set of social norms that are particularly important in the context of basic income are the norms of work, which can affect an individual's decision whether and how much to work, as well as the responsiveness of work decisions to changes in the income

[20] Using both laboratory and field experiments, various papers have demonstrated that individuals' tax compliance is affected by information about other taxpayers' compliance behaviour and their beliefs about compliance (see, for example, Alm et al., 1999, and Wenzel, 2005). This norm of compliance is often termed "tax morale" (Luttmer & Singhal, 2014).

tax system. We begin with a brief review of recent empirical evidence examining work norms before discussing the implications of such norms for the design of income taxes.

Woittiez and Kapteyn (1998) and Aronsson et al. (1999) both show that an individual's choice of hours to work depends on the hours worked by their social reference group. Consequently, a change in taxes will affect an individual's labour supply decision both directly by changing the individual's own return to work, but also indirectly by affecting the hours worked by the individual's social reference group. Using German panel data, Chadi (2014) finds evidence for a social work norm that causes disutility to unemployed people living off of public funds. He finds that this is also true of employed people receiving public funds, but that individuals are better off employed with a transfer than unemployed. Grodner and Kniesner (2006) calibrate a model where consumers maximize a Stone–Geary utility function and find evidence that labour supply decisions are affected by "spillover effects" where individuals move toward the average working hours of their reference group. Therefore, they respond to taxation in a complex way that depends on their peers' decisions. Using the 1979 National Longitudinal Survey of Youth, Weinberg et al. (2004) track individuals as they move across neighbourhoods and find that their labour market activity is correlated positively with those of their neighbours. Though this need not be due to work norms in neighbourhoods, it is one consistent explanation.

Kleven (2014) explores social norms as an explanation for the observed high levels of taxation in Norway, Sweden, and Denmark relative to the United States. The tax rate in Scandinavian countries is close to 80% when accounting for all income, payroll, and consumption taxes, but Scandinavian governments tend to use their public spending to in effect subsidize labour supply by lowering the prices of goods that are complementary to working (such as child care, preschool, and elder care). Meanwhile, the U.S. has relatively lower levels of distortionary taxes and transfers, but without the corresponding spending on goods and services that support the working population. Social norms may play an important role in determining which tax regime is observed. Kleven also demonstrates a positive correlation between the proportion of people believing that other people can be trusted and the tax-to-GDP ratio for various countries. Social beliefs about why people are poor also appear to be important. There is a weak negative correlation between the proportion of society who believe poor people are lazy and tax-to-GDP ratio. This correlation supports the idea that cultural norms and social trust are related to the level of taxation, but this relationship is not necessarily causal.

Social norms may be endogenous, in that they can evolve. Bordignon (1993) uses perceptions of fairness and ethics to endogenize the social norms of taxation by adopting "Kantian" preferences. Under Kantian preferences, individuals determine what they consider to be a fair tax burden based on their income and what they believe the average person should pay, then they choose how much to evade. Boadway et al. (2007) adapt the notion of ethical behaviour and Kantian preferences to the context of tax avoidance. Individuals first choose an ethical tax liability. If their actual tax liability is less than or equal to the ethical tax liability, then they behave ethically and choose their labour hours as if the tax rate is zero. If their actual tax liability exceeds the ethical amount, they behave "egoistically" and adjust their labour supply to reduce their tax burden, behaving as economic optimizers. This model yields higher labour supply and lower tax elasticities than models with no ethical considerations, offering an explanation for why observed labour supply and tax elasticities differ from standard economic predictions.

Incorporating norms regarding both extensive and intensive work decisions, Aronsson and Sjögren (2010) examine Pareto efficient income taxation in a two-ability-type optimal income tax model. The intensive social norm is modelled as a utility cost if hours worked deviate from some reference number of hours (either the mean or the mode of the hours worked distribution). This generates a potential wedge between an individual's marginal rate of substitution of leisure for consumption and their after-tax wage. Consequently, from the planner's perspective (who respects individual preferences), working generates an externality that can be corrected through the income tax system. The form of corrective taxation (positive or negative marginal tax rates) ultimately depends on the distribution of ability types in the population since this determines the distribution of hours worked and the reference hours. To capture the participation social norm, Aronsson and Sjögren assume the utility cost of not participating depends positively on the total number of participants in the population. The participation (extensive) social norm alone does not change the Pareto efficient income tax schedule relative to the models case without any social norms, since the optimal marginal tax rates depend only on the intensive labour supply decisions. The optimal marginal tax rates will be affected in the presence of social norms relating to hours worked (intensive-margin social norms), depending on the reference hours. In this case, the optimal tax changes to address the externalities arising from the intensive-margin norms will also affect individuals' participation decisions.

Boadway and Martineau (2016) instead focus on social norms affecting the work participation decisions using the extensive-tax model of Diamond (1980) and Saez (2002). They assume the participation decision for a type-i individual is given by

$$c_i + x(h_0, c_0) \geq c_0 + v_i \tag{42}$$

where c_i is consumption if working, c_0 is the transfer if not working, and v_i is the utility from leisure. The term $x(h_0, c_0)$ is a social norm term reflecting the moral reward of work, where h_0 is the total number of non-participants. Low (high) levels of non-participation signal that most non-participants do not have work due to luck (choice), and $x(\cdot)$ is positive (negative) and increasing (decreasing) in the transfer to non-participants. Social norms may also affect either the utility cost of not participating (which can be viewed as a form of stigma), and whether optimal participation taxes will be higher or lower will depend on the welfare weight given to non-participants and the strength of the feedback effects of the social norm on participation and government resources. Given these assumptions, multiple local stable equilibria can arise, where the optimal tax is constrained by present norms, but governments could also use the tax system to induce a change in norms. Unfortunately, the transition dynamics of shifting from one equilibrium to another (for example, from a low-participation to a high-participation economy) are complex and depend on political factors. For example, if a government initiates but does not complete a transition from one equilibrium to another for some reason (such as losing power or an electoral backlash), individuals may end up being worse off. If, however, commitment to a complete transition is possible, transitioning from one equilibrium to another has the potential to be Pareto improving. Focusing on a local equilibrium with a low (high) level of non-participation, the participation social norm will result in higher (lower) optimal participation taxes than without the social norm.

Another set of social norms that could have implications for basic income are norms of distribution governing what allocations of resources individuals feel are fair and norms of reciprocity. Kahneman et al. (1986) find that norms about fairness reduce the ability of firms to respond to market conditions in terms of making profits. For example, individuals find it unfair that a firm reduces wages in times of high unemployment. Yet if a worker left, the individual has no issue with their replacement being paid the (lower) market rate. Görges and Nosenzo (2020) in a recent handbook chapter survey the experimental evidence measuring social norms in the labour market, including the norms of reciprocity with respect to effort provision. The evidence provides some support that people believe workers should be paid some minimum reference wage even if economically inefficient, but that effort should be commensurate with wages. Low effort is unacceptable when individuals receive high wages, but acceptable when receiving significantly lower wages.

Additionally, these studies find that shirking is less acceptable when effort is observable.

In short, there are many norms that limit the idiosyncratic inequities that individuals face. These norms may not always be economically efficient. While many people worry about the effects basic income may have on labour supply, societies with strong norms of reciprocity may actually find themselves in a situation where people are mobilized to more socially useful aims. Individuals who can work in the market may continue working thanks to norms of work, while those who cannot may feel additional social impetus to "pay back" the basic income (due to norms of reciprocity) with other socially useful but economically undervalued work.

Optimal Income Taxation With Behavioural Biases

Behavioural considerations have been incorporated into optimal income tax models by Gerritsen (2016), Lockwood (2016), Bernheim and Taubinsky (2018), and Fahri and Gabaix (2019). In each case, deviations from standard behaviour take the form of biases in labour choice. To illustrate, consider the case of a linear progressive income tax in (1). Suppose the utility function takes the quasilinear form $u\big(c-h(\ell)\big)$. In the absence of bias, an individual with wage rate w solves problem (4), giving

$$h'(\ell)=(1-t)w, \quad \Rightarrow \quad \ell\big((1-t)w\big) \tag{43}$$

and the indirect utility function $v(t,b;w)$. For individuals with behavioural bias, (43) is not satisfied. Following Gerritsen (2016), define the bias by γ, where:

$$\gamma=\frac{h'(\ell)}{(1-t)w}-1. \tag{44}$$

The bias could represent various behavioural considerations discussed above, such as present bias, keeping up with the Joneses, and so on. If $\gamma>0$, the individual supplies too much labour, and vice versa. In Bernheim and Taubinsky (2018), γ varies among wage types, so we denote it γ_w.

The government treats γ_w as a deviation from optimal behaviour and gives it no weight in social welfare. The government maximizes a utilitarian social welfare function $\int_w v(t,b;w)dF(w)$ subject to an aggregate budget constraint $b=\int_w ty(w)dF(w)=t\bar{y}$, where $\bar{y}$ is average income. The problem is a straightforward extension of the standard optimal linear progressive taxation problem discussed above, and the optimal tax rate t is an extension of (8):

$$\frac{t}{1-t}=\frac{-\mathrm{Cov}\big[\beta(w),y(w)\big]}{\bar{\beta}\bar{y}\varepsilon_{\bar{y}}}+\Big(E[\gamma_w]+\mathrm{Cov}[\gamma_w,\nu(w)]\Big) \tag{45}$$

where $\varepsilon_{y(w)}$ is the elasticity of taxable income, $\beta(w)$ is the marginal utility of consumption, and $\nu(w) = \beta(w)y\varepsilon_{y(w)}/(\overline{\beta}\overline{y}\varepsilon_{\overline{y}})$ reflects how the taxpayer's marginal utility of consumption and elasticity of taxable income compare with the average. Eq. (45) differs from the standard case by the last term, which reflects the social benefit of offsetting individual biases. If $\gamma_w > 0$, the tax rate is increased to reduce the oversupply of labour. The correction will be greater to the extent that those with large biases have higher marginal utilities of consumption and elasticities of taxable income.

Bernheim and Taubinsky (2018) also consider the case where the bias is due to misperceptions about the tax rate. If individuals underestimate the income tax rate, they are less responsive to increases in the tax rate. This allows the government to set a higher tax rate and achieve more redistribution.

Lockwood (2016) assumes the behavioural distortion takes the form of present bias. Consumers put excessive weight on labour supply relative to consumption and maximize the biased utility function $u\big((1+\gamma_w)c - h(\ell)\big)$, where $\gamma_w < 0$. Their first-order condition is $h'(\ell) = (1+\gamma_w)(1-t)w$, so they undersupply labour. Assuming γ_w is the same for all, (45) reduces to:

$$\frac{t}{1-t} = \frac{-\mathrm{Cov}[\beta(w), y(w)]}{\overline{\beta}\overline{y}\varepsilon_{\overline{y}}} + \gamma.$$

The greater the present bias—the smaller is γ—the more is labour undersupplied, and the lower is the optimal marginal tax rate.

This analysis has been extended to optimal nonlinear income taxation by Gerritsen (2016) and Fahri and Gabaix (2019). The approach of Fahri and Gabaix extends that of Bernheim and Taubinsky (2018) by incorporating additional behavioural effects. Suppose we write the utility function in terms of consumption and income as $v(c, y; w) \equiv u(c, y/w)$. Maximizing this subject to the budget constraint $c = y - T(y)$ gives the first-order condition $\big(1 - T'(y)\big)v_c + v_y = 0$. Drawing on Gerritsen (2016), Fahri and Gabaix define the behavioural wedge as

$$\tau^b(y) = -\frac{\big(1 - T'(y)\big)v_c + v_y}{v_b} \tag{46}$$

where v_b is the marginal utility of income. If $\tau^b(y) > 0$, the individual overestimates the benefits of working so works too much, and vice versa. This definition captures the individual's misoptimization over labour supply discussed above.

A second behavioural effect considered in Fahri and Gabaix (2019) is a behaviour cross-influence, denoted as $\varepsilon_{y^*}(y)$. It is the compensated elasticity of earnings of an individual of income y with respect to the marginal retention rate at income level $y^* \neq y$—that is, $\big(1 - T'(y^*)\big)$. For example, if individuals mistake average for marginal tax rates, the tax rate at $y^* < y$ influences labour supply at income level y, and $\varepsilon_{y^*}(y) > 0$. Fahri and Gabaix derive the optimal marginal tax rate analogous to (17) in the presence of these two behavioural effects:

$$\frac{T'(y^*) - \tilde{\tau}^b(y^*)}{1 - T'(y^*)} = \frac{\big(1 - H(y^*)\big)\big(1 - G(y^*)\big)}{\varepsilon(y^*)y^*h(y^*)} - \int_0^\infty \frac{\varepsilon_{y^*}(y)}{\varepsilon(y^*)} \frac{T'(y^*) - \tilde{\tau}^b(y^*)}{1 - T'(y)} \frac{yh^*(y)}{y^*h^*(y^*)} dy. \qquad (47)$$

The first term on the right-hand side is the same as in the standard case of (17). The other terms capture behavioural effects. When the behavioural cross-influence elasticity $\varepsilon_{y^*}(y)$ is positive, an increase in the tax rate of individuals at income y^* causes those at income y to perceive their own tax rates as higher, leading to a decrease in the labour supply and a reduction in the optimal marginal tax rate. The tax rate is also affected by the behavioural wedge, $\tilde{\tau}^b(y)$, caused by misperceptions of the value of working. When $\tilde{\tau}^b(y) > 0$, the agent works too much, and the optimal marginal tax rate will increase.

Fahri and Gabaix (2019) note that when $\tilde{\tau}^b(y) < 0$, it is possible that $T'(y^*) < 0$. For example, individuals may underestimate the fact that more working leads to higher human capital accumulation and higher wages in the future. They argue that these biases could be particularly important at the bottom of the earnings distribution so that the marginal tax is negative at low incomes. This would provide a behavioural rationale for an EITC different than that proposed by Saez (2002). Evidence from the Self-Sufficiency Project in British Columbia and New Brunswick in the 1990s does not support this argument for participation subsidies. Part of the purpose of these experiments was to evaluate the idea that providing an incentive to work would cause wages to rise and improve employment prospects. As noted above, Card and Hyslop (2005) found that while employment was increased while workers were supported by the experiment, once it concluded, the employment effects disappeared. There was apparently little human capital generation for low-income experiment participants.

Gerritsen (2016) derives the analogue of (47) without the behavioural cross-influence, so $\varepsilon_{y^*}(y) = 0$. He empirically estimates the behaviour wedge $\tau^b(y)$, which is a sufficient statistic for the degree of individual misoptimization. He uses information on life satisfaction from the British Household Panel Survey as a measure of individual well-being based on respondents' answers to the question whether they would prefer to work fewer, greater, or the same number of hours at their current hourly wage rate. He finds support for a positive relationship between a worker's current income and the amount they are working relative to their well-being maximizing amount. Consequently, the marginal tax rates on low-income individuals should be lower to induce greater labour supply, lending some empirical support to the assertion of Fahri and Gabaix (2019).

Lessons From Optimal Income Taxation for Basic Income Policy

A basic income guarantee is a common feature of the optimal income tax literature. The standard approach in the literature is to choose an income tax structure that maximizes a social welfare function subject to a government budget constraint and taking into account individual behavioural responses to taxation. The social welfare function satisfies non-negative aversion to inequality in utility, which leads to aversion to income inequality. Underlying the budget constraint is an assumed knowledge of the characteristics of the population, such as their distribution of skills, as well as the economy's production technology. Restrictions may be imposed on the form of the income tax structure—for example, linear progressive, piecewise linear, or nonlinear. The government may also face informational constraints. Individual characteristics or effort determining earned income may be unobservable. For individuals who are unemployed, the government may not know whether they are voluntarily or involuntarily unemployed, or whether they are able to work. Even if the government can observe whether an unemployed person chooses not to work, is unable to work, or is able to work but cannot find a job, it still has to decide what sizes of transfers to give to each class of unemployed, and that will involve some normative judgment. The exact structure of the optimal tax-transfer system and the size of the basic income guarantee will depend upon all of these factors.

The models used to analyze optimal income taxation are very diverse. The results are model-specific and depend critically on assumed social objectives. Nonetheless, there are some common themes running through the literature. We focus on the results that are most important for basic income. Two main dimensions are relevant. One is the size of the basic income guarantee and whether it differs for different types of persons. The size of the basic income guarantee can be interpreted as the transfer to those with no taxable income—that is, the intercept of the optimal income tax schedule along the vertical axis. The second is the structure of income tax rates near the bottom of the income distribution, including both marginal tax rates and participation tax rates. Of particular interest is whether marginal tax rates are relatively high at the bottom compared with at higher income levels, and whether the participation tax rate is positive or negative.

Our purpose in this section is to summarize the key results recounted in more detail above. The starting point is the Mirrlees optimal income tax model based on intensive-margin labour supply decisions. Individuals have the same preferences but different fixed wage rates and face a perfectly elastic demand for labour. The optimal income tax structure features marginal tax rates between zero and unity, but the pattern of marginal tax rates is ambiguous. Simulations and analytical approaches suggest that marginal tax rates are initially declining and then increase moderately, except under a maximin social welfare function where they tend to fall monotonically. Average tax rates tend to rise over most of the income distribution, although they may peak before the top. The size of basic income is generally hard to characterize. It tends to be higher the more averse to inequality is social welfare. It is also higher the greater is the general value-added tax rate.

If the government is restricted to a linear progressive income tax, the fixed marginal tax rate rises with the basic income guarantee. The basic income is higher the greater is the aversion to inequality and lower the larger are government expenditure requirements. If the income tax is piecewise linear, marginal tax rates can either rise or fall starting from the lowest tax bracket. Under reasonable assumptions, the marginal tax rate is higher at the bottom than higher up, and basic income is higher than in the linear progressive tax case.

Policy prescriptions can change when individual labour supply decisions are restricted to be discrete—that is, along extensive margins. In the simplest case of Diamond (1980) and Saez (2002), individuals choose only whether to participate in the labour market. Those who participate are restricted to a job suitable to their skills and paying job-specific earnings. The tax system consists of discrete taxes based on earnings, including a transfer (basic income) to those with zero earnings. An innovative finding in this pure participation model is that if an increment of income to non-participants is given more social weight than an increment to participants and if aversion to inequality is not too high, a participation subsidy to the lowest-skilled participants is optimal. This is of practical relevance since it corresponds with the CWB in Canada and the EITC in the U.S. While this result has generated much attention, it is restrictive. It does not apply if the social weight given to non-participants—who include individuals of various skills who choose not to work—is less than the social weight given to low earners. Nor does it apply if the social welfare function exhibits a high degree of aversion to inequality, such as the maximin social welfare function. More generally, the pattern of participation tax rates is ambiguous in the extensive-margin model. If the elasticity of participation is constant or increasing with income, the participation tax rate will be increasing.

Adding other discrete choices reduces the likelihood of participation subsidies. If, as in Saez (2002), individuals can choose both whether to participate and, if so, whether to take the job most suited to their skills or the one that is one skill level lower, the case for a participation subsidy at the lowest earnings level is reduced. A participation subsidy will be optimal only if the elasticity of participation is large enough relative to the elasticity of job choice.

Adding more margins of decision-making tempers the case for participation subsidies further. If individuals can make both an extensive-margin choice—whether to participate—and an intensive-margin choice—how much to work—the chances that a participation subsidy is optimal is reduced. As well, marginal tax rates are reduced, and the optimal income tax is overall less progressive. Generally, the more decisions are made, the higher is the efficiency cost of taxation, the lower marginal tax rates should be, and the less likely is it that a participation subsidy is optimal.

Some empirical evidence informs the analytical approach. The efficiency cost of income taxes is affected by the elasticity of taxable income. Estimates suggest that this increases with income levels. However, there is little evidence about the size of the elasticity at the bottom, nor about the income effect of transfers at the bottom. This makes it difficult to evaluate the incentive effects of a basic income. There is compelling evidence that the elasticity of participation is very low in the United States, which would suggest that the efficiency cost of participation subsidies are low.

Empirical evidence also lends some support to the social welfare maximization approach and its consequences for basic income. Estimates of social preferences based on existing tax-transfer systems suggest that social weights are highest for those who are not working. They are lower for the working poor and tend to fall with income. There are differences among countries. The highest groups have higher social weights in the U.S. relative to other countries, especially Sweden, where social weights are negative at the top. Having a study for Canada would be useful. The results from the stated-preference approach to estimating social preferences using surveys are more agnostic. Some studies refute the assumption that redistribution should be based on social welfare maximization alone. They tend to show that citizens give comparable weight to equal sacrifice as an objective, which implies that some weight should be given to where people start when evaluating optimal taxation. Other studies show that those surveyed place weight on the process by which redistribution occurs—for example, on progressivity, per se, independent of its effect on actual income distributions.

Other aspects of the standard optimal income tax model have been questioned. The standard model assumes that social welfare depends on individual utilities, which in turn depend on consumption and labour supplied by the individuals. There has been a sizable literature investigating the consequences on interdependent utility functions for redistribution. An important example of this is the tax treatment of voluntary transfers, such as bequests or charitable donations, which are presumed to be motivated by altruism. Two issues are relevant. One is whether altruistic benefits should count for social welfare purposes. That is, when a voluntary transfer generates benefit both for the donor through altruism and for the recipient as a gain in income, should both be counted as social welfare. Adherents to an individualistic view of social welfare argue that altruistic benefits should count given that they are reflected in revealed preferences. Those opposed argue that such double-counting is not warranted for reasons we have outlined above.

The second issue is what are the implications for optimal income taxation if interdependent utilities fully count for social welfare purposes. Altruism may be invoked as the rationale for redistribution instead of a social welfare function. In this case, altruism of the rich counts and supersedes the utility of the poor. The optimal amount of redistribution is determined by the amount that the rich would willingly give, and in what form. If the rich have paternalistic altruism, they may prefer to give conditional transfers or in-kind transfers instead of unconditional transfers. Because there is a free-rider problem among the rich, government redistribution on behalf of the rich internalizes the free-riding problem. This setting would generate less redistribution than the social welfare approach assuming the marginal utility of income of the rich exceeds the marginal value they put on transfers to the poor. The altruistic outcome is inconsistent with social welfare maximization based on individual preferences, and it does not accord with the amount of redistribution observed in practice or with social weights estimated as above.

Optimal income taxation can take interdependent utility functions into account if desired. The special case where individuals obtain utility from status goods or "keeping up with the Joneses" calls for greater progressivity of the income tax on efficiency grounds, thereby supplementing equity arguments for redistribution. In the more general case where utility functions have an argument reflecting envy of other individuals' consumption, optimal marginal tax rates are increased. Conversely, when individuals are altruistic so benefit from other individuals' consumption, optimal marginal tax rates are correspondingly lower, and can even be negative for some taxpayers. If the altruism applies particularly to low-income households, the size of the optimal basic income is increased.

A limiting feature of the standard approaches to optimal income taxation is that the demand for labour of all types is perfectly elastic, so any changes in labour supply automatically lead to changes in employment at given wage rates. This is unrealistic for two reasons: wages rates are endogenous, and full employment may not be guaranteed. The optimal income tax literature has been revised to address both issues.

Suppose first that the demand for labour is not perfectly elastic, but labour markets are competitive and wage rates are endogenous. A decrease in the marginal income tax rate at the top will increase the ratio of high-wage to low-wage labour and cause the relative wage of the low-skilled to rise. An increase in the marginal income tax rate at the bottom will have a similar effect, so overall, a steepening of the rate structure will improve equity. A higher participation tax and a higher basic income will have the same effect by reducing low-wage labour supply. More generally, a negative income tax will achieve more redistribution than a participation subsidy (CWB) accompanied by income tax financing. It has been estimated that a large proportion of the EITC in the U.S. accrues to employers via wage rate reductions. These effects will be offset to the extent that labour market experience increases skills. In this case, policies that encourage low-income labour supply will increase low-income wages, although empirical studies based on experimental projects do not support this experience-enhancing effect. Finally, optimal tax progressivity can also fall with social mobility.

Next, suppose some individuals are unemployed or underemployed either because they are unable to work full-time or because they are unable to find a job. If those unable to work—for example, the disabled—can be identified, they may be offered a basic income that is different from those who choose not to work. Typically, they can only be imperfectly identified, or "tagged," so some disabled are not tagged and vice versa. The government can offer separate tax-transfer systems to tagged (lower-income) and untagged (higher-income) groups, which differ in progressivity. In addition, funds can be transferred implicitly from untagged to tagged groups as a whole, making the system more progressive overall. Tagging enables the government to make higher transfers to the disabled who are tagged, but at the cost of the untagged disabled being made worse off. The more averse to inequality is the government, the less valuable tagging will be. In the maximin case, tagging would not be used as long as some disabled are not tagged. In principle, tagging could be used to differentiate the tax-transfer system applicable to identifiable sub-groups in the population, such as gender, family size, age, or province of residence.

The prospect of tagging raises the general issue of how to treat persons with different work habits. Social weights may be relatively high for those who work very little because of disability. However, lower social weight may be attached to those for whom the intensity of work reflects different preferences for leisure. In this case, the government may want to redistribute from the lowest-income persons who have the highest preference for leisure to those with slightly higher income who have the same skills but are more industrious. That is, the least well-off persons are not necessarily the lowest-income ones. This leads to negative marginal tax rates at the bottom, similar to participation subsidies. A practical problem is that the lowest income groups may contain both those with disabilities and those with high preferences for leisure, and that will tend to make social weights attributed to the lowest- and zero-income individuals higher.

Suppose now that some persons are involuntarily unemployed. We can distinguish between long-term and short-term involuntary unemployment. The former case includes persons who have not been able to find a job, as well as those who are structurally unemployed and whose skills are obsolete. Transfers to the long-term unemployed are redistributive in nature and may be delivered through stand-alone programs such as welfare or through the income tax system as a basic income. The optimal income tax literature addresses long-term involuntary unemployment using a model in which the unemployment is due to frictions in the labour market, typically search or matching frictions. The wage rate is determined by bargaining. The tax system directly affects the equilibrium wage, which in turn affects individuals' search/participation/labour supply decisions. Optimal marginal tax rates are positive and higher than in the full employment model, and the average tax rate is increasing in income. Optimal participation taxes are generally lower than with full employment and may be negative unless aversion to inequality is high (e.g., maximin). Basic income or transfers to the unemployed will generally be greater than in models without involuntary unemployment, assuming the government cannot distinguish between the voluntary and involuntary unemployed.

Underlying these results is the process of wage determination through bargaining. Alternative approaches to involuntary unemployment have different policy implications. If unemployment is the result of efficiency wages—that is, wages that are above the market-clearing level to induce worker effort—payments to the involuntary unemployed can be counterproductive. A higher basic income transfer would make it more expensive for firms to induce worker effort, so wages, and therefore unemployment, would rise. Minimum wages that induce some unemployment might also be used as a policy instrument alongside a basic income. If participation subsidies at low incomes are optimal, it may be welfare-improving to impose a binding minimum wage. Inducing some involuntary unemployment of low-skill workers saves tax revenue, since transfers to low-income workers are higher than to the unemployed, and increases the income of the low skilled who are still working, which together can outweigh the efficiency costs of the induced unemployment and improve welfare.

If involuntary unemployment is short term, reflecting temporary layoffs or frictional unemployment, transfers to the unemployed fulfill an insurance purpose as well. Workers may be unable to self-insure, so public unemployment insurance might be efficiency-enhancing. At the same time, if the inability to self-insure mainly affects lower-income workers, unemployment insurance may well have an equity objective as well. As with redistributive taxation, unemployment insurance can also have behavioural consequences since it can affect the incentive to search and to participate in the labour market. Unemployment insurance can also correct for inefficient wage bargaining. There is limited literature examining unemployment insurance in a redistributive setting. That which exists suggests that if the government can distinguish the voluntary from the involuntary unemployed, it is optimal to target redistribution using the income tax system, and to use unemployment insurance to insure workers against job loss. The income tax system would be constrained by labour market behaviour, including wage determination, while the unemployment insurance system would trade off insurance versus search inefficiency. If the government could not distinguish the voluntary from the involuntary employed, these targeting prescriptions might not be clear-cut. Further analysis of unemployment insurance design in an optimal income tax setting would be useful. Optimal income taxation can also account for other behaviour considerations arising from the presence of social norms or biases in decision-making. With social norms, individuals' utility depends not only on their own work decisions, but also on the work decisions of others. Consequently, norms of work are similar to externalities that the income tax system can be used to correct for. If the work norm is about hours worked, then it is the marginal tax rate that will be adjusted and may be higher or lower than without the work norm depending on the skill distribution. If it is a work participation norm, then it is the participation tax rate that will be adjusted relative to without the work norm. With biases in labour supply decisions, there will be a wedge between the after-tax return to working and the marginal utility cost of working, which can be corrected for by the income tax system and may result in a more or less progressive tax system than without the bias, depending on the form of the bias (misperception of tax rates, present bias). Individuals may also underestimate that working more leads to higher future wages, and, consequently, negative marginal tax rates may be optimal, although this wage-enhancing effect is not supported by empirical evidence.

To summarize, the optimal income tax literature suggests a basic income integrated with the income tax system with a marginal tax rate (or taxback rate) starting at $b = 0$ that exceeds that of the following income tax brackets. A different basic amount could be given to those who are disabled.

The case for a participation subsidy at low earnings levels is ambiguous at best. If the voluntary unemployed could be identified, a case could be made for their basic income being less than those unable to find a job. As mentioned earlier and discussed further below, the basic income could be operated as a separate component of the income tax system as a refundable tax credit. Doing so would enhance the flexibility to adjust transfers within the tax year. The basic income system would operate separately from unemployment insurance and would supersede existing welfare systems. It would be unconditional and would complement other social services. Basic income would apply to those who are temporarily unemployed. It would be based partly on insurance principles but could incorporate some preferential treatment of low-income workers who are least able to self-insure. The unemployment insurance system, unlike basic income, would require monitoring to ensure workers are eligible to receive, where eligibility could be restricted to those who lost their jobs through no fault of their own and who are willing to take suitable jobs that are offered. In the event that workers are unable to find another job, they would move to the basic income system.

Implications for Designing a Basic Income Guarantee in Canada

We have stressed that the optimal income tax literature prescribes a basic income guarantee as an implicit component of the optimal tax structure. However, little attention is paid to practical tax design and institutional considerations. In this final section, we briefly review some issues that would arise in implementing a basic income guarantee through the income tax system in the Canadian federal setting. A basic income delivered through the income tax system would be administered by the relevant tax administration, such as the CRA and/or one or more of the provincial tax administrations. Like income taxes, transfer entitlements would be determined on a self-reporting basis, with auditing, transfer adjustment, and possibly penalization as enforcement devices. This is in contrast to provincial welfare systems or the federal EI system where eligibility is determined by ex ante application and verification, and by ongoing monitoring to ensure compliance. This gatekeeping approach is the source of concern over stigmatization that is largely avoided by self-reporting, but which also has some advantages in terms of flexibility to respond to changes in circumstances that are discussed further below. Deficiencies in take-up can be a concern in both systems: reluctance to apply in the case of welfare, and failure to file a tax return in the case of transfers delivered through the tax system.

Recall that there are three equivalent ways to implement a basic income guarantee based on an optimal income tax system, in which all individuals face the same nonlinear income tax schedule with an intercept of b at zero income. First, a universal basic income b could be paid up front to all tax-filers, and then a prescribed tax schedule would be applied to all income, including b. Second, an income-contingent basic income could be delivered as a transfer of negative tax liabilities calculated after taxes have been assessed. Finally, a basic income transfer could be delivered separately as a refundable tax credit analogous to existing refundable tax credits (e.g., Canada Child Benefit, Working Income Tax Benefit) and to Old Age Security/Guaranteed Income Supplement (OAS/GIS). All three can be designed to give the same final outcome. They differ mainly in timing. The universal basic income involves an upfront transfer of b to all individuals and collects taxes to finance it ex post. The latter two pay only negative taxes as refunds to those entitled to them. All three are administered through the income tax system using self-reporting and possibly ex post auditing. This requires that all basic income recipients file income tax returns, which may require some public assistance. The refundable tax credit version makes it easier to disburse payments on a monthly or quarterly basis and to respond to changes in individuals' incomes during the tax year. We assume this version in what follows.

The basic income credit involves the designation of the basic income amount b and the taxback rate or rate schedule. The value of b could differ according to personal circumstances. A smaller b could apply to children and be paid on their behalf to parents. A smaller b could also be paid to those in multi-adult families—for example, using the square-root equivalence-scale rule. Total payments for a family of n adults would be $\sqrt{n} \cdot b$, suitably taxed back and paid in equal shares to all adults. The taxback could be based on family income. A larger value of b could be paid to those judged to be disabled. The basic income amount could be based on the poverty rate, such as the Market Basket Measure (MBM) adopted by the Canadian government (Employment and Social Development Canada, 2018). The MBM varies by location, and allowing b to vary by location would complicate the basic income guarantee. An option would be for the federal government to choose a common national value of b but allow the provinces to choose province-specific values as discussed below. Note that the MBM takes account of the need for market-based purchases of goods and services. To the extent that the government provides goods and services to households either universally or based on need, that would in principle reduce the MBM and therefore the value of b. If the basic income replaced provincial social assistance systems, the delivery of needs-based in-kind transfers would have to be addressed, since eligibility for the latter is currently often determined by social assistance eligibility.

The choice of the taxback rate would be influenced by incentive effects such as elasticities of taxable income and income effects to the extent that reliable estimates exist. As in the optimal tax literature, effective marginal tax rates at the bottom could exceed those higher up, especially since the value of b is likely to be much higher than in the current system. It may be desirable for the basic income credit to include a participation incentive analogous to the CWB, as in Koebel and Pohler (2019). More information about labour market responses to transfers to low-income persons would be helpful, including effects on both labour supply and demand. The basic income credit should be harmonized with other tax credits in the income tax system to avoid the stacking of taxback rates. For example, the value of all credits could be aggregated and a common taxback rate applied, analogous to the Universal Credit system in the U.K.

The design of a basic income for Canada must take the federalism dimension into account. Both federal and provincial governments redistribute, and most harmonize their income tax systems and take advantage of CRA tax administration. Existing income-conditioned transfers include provincial welfare and disability assistance, federal OAS/GIS, and various federal and provincial refundable and nonrefundable tax credits. A move to a national basic income guarantee delivered through the tax system would ideally involve federal and provincial collaboration, both to retain distinct federal and provincial redistributive roles and to continue to take advantage of the common tax administration. One option is to exploit the income tax collection system to design a basic income guarantee using the refundable tax credit approach, as proposed by Boadway et al. (2018b). The federal government would take the initiative by implementing a federal basic income b_F. Each province could then choose whether to implement its own basic income b_{Pi}. These choices could together reflect the MBM poverty rates. A common taxback rate would apply to the aggregate of federal and provincial choices of $b_F + b_{Pi}$. The national basic income guarantee level would vary among provinces, reflecting provincial preferences and needs.

This two-stage procedure would allow the basic income guarantee to be phased in and would respect the joint federal and provincial interests in redistribution. It would, however, rely on the federal government to initiate the process. In principle, the initiation could begin with one or more provinces, but some federal collaboration would be required. A single province implementing a basic income guarantee would have to secure agreement with the federal government to revise the relevant federal–provincial income tax collection agreement to allow for the basic income refundable tax credit. Once in place, other provinces and/or the federal government could follow suit.

An important issue in designing a basic income guarantee is how to finance it. One possibility has been explored by Stevens and Simpson (2017), Boadway et al. (2018b), and Koebel and Pohler (2019). They propose to reallocate revenues from provincial social assistance and most existing refundable and nonrefundable tax credits to a basic income delivered as a refundable tax credit. The relative size of provincial and federal components would reflect the amounts of federal and provincial funds currently allocated to their respective transfer programs. Other financing options include eliminating some other existing tax expenditures, using carbon tax revenues, and increasing income or sales tax rates. Most proposals envision keeping social insurance programs in place, along with social services and in-kind transfers. As mentioned, the principle would be to relate the basic income guarantee to the MBM poverty measure so would be intended to ensure that all individuals are able to afford to purchase bundles of goods on the market comparable to those included in the MBM. Goods and services provided by the government would not affect the basic income guarantee except to the extent that they affect the MBM.

Other important design issues would have to be addressed, some of which have been touched on in the optimal income tax literature. One is the treatment of working versus non-working individuals. The latter can include those who choose not to work—the voluntary unemployed—and those who are unable to work or find a job—the involuntary unemployed. The treatment of the voluntary unemployed relative to those working involves a contentious value judgment: Should society punish those who choose not to work? Those who subscribe to the equality of opportunity approach argue that individuals of equal productivity should take responsibility for their decision and should be neither rewarded nor punished for exercising their preferences. Others argue that society is unlikely to condone giving substantial transfers to persons who choose to be idle. Some argue that a basic income guarantee should be conditioned on participation in a broad sense, including, for example, working, volunteering, caring, or training. Optimal income tax analysis offers some support to participation subsidies, such as the CWB, and, as mentioned, recent studies have found that the empirical effect of participation subsidies is ambiguous and possibly negligible.

Transfers to the involuntary unemployed might ideally differ from those voluntarily unemployed. Those who are unable to work because of disability might warrant a higher transfer, if only because it is more costly for them to enjoy a minimum standard of living. This requires that they be identified, and that may be difficult to do accurately. To the extent that the disabled cannot be differentiated from the voluntary unemployed, a common basic income would have to apply to all, and that may both work to the advantage of the voluntary unemployed and enhance the case for a participation subsidy. Similarly, if those who are involuntarily unemployed because they are unable to find a job can be identified, they may be afforded a higher transfer. Identifying them may not be easy, unless one is prepared to engage in significant monitoring. Those who are temporarily unemployed may be eligible for unemployment insurance distinct from basic income.

The challenges of differentiating among the voluntary unemployed, the disabled, and the involuntary unemployed involve only one dimension of information problems a government will face in implementing a basic income guarantee. Another equally important problem is knowing an individual's income on which to base a transfer. There are two practical instances of this. One concerns income volatility and the difficulties it implies for delivering transfers through the income tax system. Such transfers ought to be based on current income, but given the time between income tax filing and final reconciliation, there can be significant lags in updating income information. Persons whose income falls during the tax year may have to wait a long time to receive the basic income transfer to which they are entitled. In the meantime, they may face significant liquidity problems that prevent them from purchasing necessary goods and services. At the same time, those whose income goes up will be paid more than they are entitled to, and it will be hard to recoup that from them if they are living on low incomes. This problem might be addressed by instituting a separate mechanism outside of the tax system for dealing with income changes, but this would have to be designed in such a way as to avoid stigmatization of heavy compliance costs. The system could involve an extended EI system to include precarious workers, the self-employed, and gig workers. Transfers to them could react reasonably quickly to income changes, and they could be reconciled ex post with basic income delivered through the tax system.

A second information problem results from the possibility that basic income recipients misreport their income. They may simply under-report when taxes are filed to reduce the taxback of the basic income transfer, or they may engage in non-market or underground economy activities and fail to report income earned there. The problem is mitigated to the extent that third-party reporting of income applies, but for the self-employed, income concealment is harder to detect. This problem already exists in the welfare system, but there is relatively little empirical evidence of its extent. The problem of income misreporting is related to societal norms. It is possible that social norms will improve with a basic income guarantee that eliminates the stigma of welfare programs, but that is by no means certain.

References

Akerlof, G. (1978). The economics of "tagging" as applied to the optimal income tax, welfare programs, and manpower training. *American Economic Review*, *68*, 8–19.

Ales, L., Kurnaz, M., & Sleet, C. (2015). Technical change, wage inequality, and taxes. *American Economic Review*, *105*, 3061–3101.

Alesina, A., & Giuliano, P. (2011). Preferences for redistribution. In A. Bisin & J. Benhabib (Eds.), *Handbook of social economics* (pp. 93–132). North-Holland Publishing Company.

Alesina, A., Ichino, A., & Karabarbounis, L. (2011). Gender-based taxation and the division of family chores. *American Economic Journal: Economic Policy*, *3*, 1–40.

Alesina, A., Miano, A., & Stantcheva, S. (2019). *Immigration and redistribution* (NBER Working Paper No. 24733).

Alesina, A., Stantcheva, S., & Teso, E. (2018). Intergenerational mobility and preferences for redistribution. *American Economic Review*, *108*, 521–554.

Alm, J., McClelland, G. H., & Schulze W. D. (1999). Changing the social norm of tax compliance by voting. *Kyklos*, *52*, 141–171.

Apps, P., & Rees, R. (2018). Optimal family taxation and income inequality. *International Tax and Public Finance*, *25*, 1093–1128.

Apps, P., Long, N. V., & Rees, R. (2014). Optimal piecewise linear income taxation. *Journal of Public Economic Theory*, *16*(4), 523–545.

Aronsson, T., & Sjögren, T. (2010). Optimal income taxation and social norms in the labor market. *International Tax and Public Finance*, *17*, 67–89.

Aronsson, T., Blomquist, S., & Sacklén, H. (1999). Identifying interdependent behaviour in an empirical model of labour supply. *Journal of Applied Econometrics*, *14*, 607–626.

Arrow, K. J. (1951). *Social choice and individual values*. Wiley.

Atkinson, A. (2015). *Inequality: What can be done?* Harvard University Press.

Atkinson, A., & Stiglitz, J. (1980). *Lectures on public economics*. McGraw-Hill.

Banks, J., & Diamond, P. (2010). The base for direct taxation. In S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, & J. Poterba (Eds.), *Dimensions of tax design: The Mirrlees Review* (pp. 548–648). Oxford University Press.

Bargain, O., Dolls, M., Neumann, D., Peichl, A., & Siegloch, S. (2014a). Comparing inequality aversion across countries when labor supply responses differ. *International Tax and Public Finance*, *21*, 845–873.

Bargain, O., Dolls, M., Neumann, D., Peichl, A., & Siegloch, S. (2014b). Tax-benefit revealed social preferences in Europe and the US. *Annals of Economics and Statistics*, *113/114*, 257–289.

Bastani, S., & Lundberg, J. (2017). Political preferences for redistribution in Sweden. *The Journal of Economic Inequality*, *15*, 345–367.

Bastani, S., Blomquist, S., & Micheletto, L. (2019). Nonlinear and piecewise linear income taxation, and the subsidization of work-related goods. *International Tax and Public Finance*, *26*, 806–834.

Bernheim, B. D., DellaVigna, S., & Laibson, D. (Eds.). (2018). *Handbook of behavioral economics* (Vol. 1). North-Holland Publishing Company.

Bernheim, B. D., DellaVigna, S., & Laibson, D. (Eds.). (2019). *Handbook of behavioral economics* (Vol. 2). North-Holland Publishing Company.

Bernheim, B. D., & Taubinsky, D. (2018). Behavioral public economics. In B. D. Bernheim, S. DellaVigna & D. Laibson (Eds.), *Handbook of behavioral economics* (Vol. 1, pp. 381–516). North-Holland Publishing Company.

Beaudry, P., Blackorby, C., & Szalay, D. (2009). Taxes and employment subsidies in optimal redistribution programs. *American Economic Review*, *99*(1), 216–242.

Blackorby, C., & Donaldson, D. (1988). Cash versus kind, self selection and efficient transfers. *American Economic Review*, *78*, 691–700.

Blumkin, T., & Sadka, E. (2005). Income taxation with intergenerational mobility: Can higher inequality lead to less progression? *European Economic Review*, *49*, 1915–1925.

Boadway, R. (2012). *From optimal tax theory to tax policy: Retrospective and prospective views*. MIT Press.

Boadway, R., & Cuff, K. (1999). Monitoring job search as an instrument for targeting transfers. *International Tax and Public Finance*, *6*, 317–337.

Boadway, R., & Cuff, K. (2001). A minimum wage can be welfare-improving and employment-enhancing. *European Economic Review*, *45*, 553–576.

Boadway, R., & Cuff, K. (2014). Monitoring and optimal income taxation with involuntary unemployment. *Annals of Economics and Statistics*, *113/114*, 121–157.

Boadway, R., & Cuff, K. (2015). Tax treatment of bequests when donor benefits are discounted. *International Tax and Public Finance*, *22*, 604–634.

Boadway, R., & Cuff, K. (2018). Optimal unemployment insurance and redistribution. *Journal of*

Public Economic Theory, *20*, 303–324.

Boadway, R., Cuff, K., & Koebel, K. (2018a). Can self-financing redeem the basic income guarantee? Disincentives, efficiency costs, tax burdens, and attitudes: A rejoinder. *Canadian Public Policy*, *44*, 447–457.

Boadway, R., Cuff, K., & Koebel, K. (2018b). Designing a basic income guarantee for Canada. In E. Goodyear-Grant, R. Johnston, W. Kymlicka, & J. Myles (Eds.), *Federalism and the welfare state in a multicultural world* (pp. 101–129). McGill-Queen's University Press.

Boadway, R., Cuff, K., & Marceau, N. (2008). Design of assistance programs to address real income volatility. In D. Jolliffe & J. P. Ziliak (Eds.), *Income volatility and food assistance in the United States* (pp. 217–258). W. E. Upjohn Institute for Employment Research.

Boadway, R., Cuff, K., & Marchand, M. (2000). Optimal income taxation with quasi-linear preferences revisited. *Journal of Public Economic Theory*, *2*, 435–460.

Boadway, R., & Jacquet, L. (2008). Optimal marginal and average income tax rates under maximin. *Journal of Economic Theory*, *143*, 425–441.

Boadway, R., Marceau, N., & Mongrain, S. (2007). Redistributive taxation under ethical behaviour. *Scandinavian Journal of Economics*, *109*, 505–529.

Boadway, R., Marceau, N., & Sato, M. (1999). Agency and the design of welfare systems. *Journal of Public Economics*, *73*, 1–30.

Boadway, R., & Martineau, N.-G. (2016). Optimal redistribution with endogenous social norms. *The Scandinavian Journal of Economics*, *118*(3), 524–556.

Boadway, R., & Pestieau, P. (2006). Tagging and redistributive taxation. *Annales d'Économie et de Statistique*, *83/84*, 2–25.

Boadway, R., & Song, Z. (2016). Indirect taxes for redistribution: Should necessity goods be favored? *Research in Economics*, *70*, 64–88.

Boone, J., & Lans Bovenberg, A. (2013). Optimal taxation and welfare benefits with monitoring of job-search. *International Tax and Public Finance*, *20*, 268–292.

Bordignon, M. (1993). A fairness approach to income tax evasion. *Journal of Public Economics*, *52*, 345–362.

Bosch, N., Jongen, E., Leenders, W., & Möhlmann, J. (2019). Non-bunching at kinks and notches in cash transfers in the Netherlands. *International Tax and Public Finance*, *26*, 1329–1352.

Bourguignon, F., & Spadaro, A. (2012). Tax-benefit revealed social preferences. *The Journal of Economic Inequality*, *10*, 75–108.

Card, D., & Hyslop, D. R. (2005). Estimating the effects of a time-limited earnings subsidy for welfare-leavers. *Econometrica*, *73*, 1723–1770.

Chadi, A. (2014). Regional unemployment and norm-induced effects on life satisfaction. *Empirical Economics*, *46*, 1111–1141.

Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy*, *116*, 173–234.

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

Chetty, R., & Looney, A. (2006). Consumption smoothing and the welfare consequences of social insurance in developing economies. *Journal of Public Economics*, *90*, 2351–2356.

Chetty, R., & Saez, E. (2013). Teaching the tax code: Earnings responses to an experiment with EITC recipients. *American Economic Journal: Applied Economics*, *5*, 1–31.

Choné, P., & Laroque, G. (2005). Optimal incentives for labor force participation. *Journal of Public Economics*, *89*, 395–425.

Choné, P., & Laroque, G. (2010). Negative marginal tax rates and heterogeneity. *American Economic Review*, *100*(5), 2532–2547.

Choné, P., & Laroque, G. (2011). Optimal taxation in the extensive model. *Journal of Economic Theory*, *146*, 425–453.

Christiansen, V. (2015). Optimal participation taxes. *Economica*, *82*, 595–612.

Christiansen, V., & Jansen, E. S. (1978). Implicit social preferences in the Norwegian system of indirect taxation. *Journal of Public Economics*, *10*, 217–245.

Cole, G. D. H. (1935). *Principles of economic planning*. Macmillan.

Congdon, W. J., Kling, J., & Mullainathan, S. (2011). *Policy and choice: Public finance through the lens of behavioral economics*. Brookings Institution Press.

Cremer, H., Gahvari, F., & Lozachmeur, J.-M. (2010). Tagging and income taxation: Theory and an application. *American Economic Journal: Economic Policy*, *2*, 31–50.

Cremer, H., Marchand, M., & Pestieau, P. (1996). The optimal level of unemployment insurance benefits in a model of employment mismatch. *Labour Economics*, *2*, 407–420.

Cuff, K. (2000). Optimality of workfare with heterogeneous preferences. *Canadian Journal of Economics. Revue Canadienne d'Economique*, *33*, 149–174.

de Caritat, Marie-Jean-Antoine-Nicolas, marquis de Condorcet. (1988). *Esquisse d'un tableau* *historique des progres de l'esprit humain* (1st ed.). Groupe Flammarion. (Original work published 1795)

Deaton, A. (1979). Optimally uniform commodity taxes. *Economics Letters*, *2*, 357–361.

Diamond, P. A. (1980). Income taxation with fixed hours of work. *Journal of Public Economics*, *13*, 101–110.

Diamond, P. A. (1998). Optimal income taxation: An example with a U-shaped pattern of optimal marginal tax rates. *American Economic Review*, *88*(1), 83–95.

Diamond, P. A. (2006). Optimal tax treatment of private contributions for public goods with and without warm glow preferences. *Journal of Public Economics*, *90*(4–5), 897–919.

Diamond, P. A., & Saez, E. (2011). The case for a progressive tax: From basic research to policy recommendations. *Journal of Economic Perspectives*, *25*, 165–190.

Ebert, U. (1992). A reexamination of the optimal nonlinear income tax. *Journal of Public Economics*, *49*, 47–93.

Employment and Social Development Canada. (2018). *Opportunity for all: Canada's first poverty reduction strategy*. Employment and Social Development Canada.

Elster, J. (1989). Social norms and economic theory. *Journal of Economic Perspectives*, *3*, 99–117.

Farhi, E., & Gabaix, X. (2020). Optimal taxation with behavioral agents. *American Economic Review*, *110*, 298–336.

Fleurbaey, M., & Maniquet, F. (2011). *A theory of fairness and social welfare*. Cambridge University Press.

Fong, C. (2001). Social preferences, self-interest and the demand for redistribution. *Journal of Public Economics*, *82*, 225–246.

Fourier, C. (1967). *La fausse industrie* [The false industry]. Anthropos. (Original work published 1836)

Frank, R. H. (2005). Positional externalities cause large and preventable welfare losses. *American Economic Review*, *95*, 137–141.

Friedman, M. (1962). *Capitalism and freedom*. University of Chicago Press.

Galbraith, J. K. (1986). *A view from the stands: Of people, politics, military power, and the arts*. Houghton Mifflin.

Gans, J., & Smart, M. (1996). Majority voting with single crossing preferences. *Journal of Public Economics*, *59*, 219–237.

George, H. (1871). *Our land and land policy, nation and state*. White and Bauer.

Gerritsen, A. (2016). Optimal taxation when people do not maximize well-being. *Journal of Public Economics*, *144*, 122–139.

Ghatak, M., & Maniquet, F. (2019). Universal basic income: Some theoretical aspects. *Annual Review of Economics*, *11*, 895–928.

Gordon, J. (1989). Individual morality and reputation costs as deterrents to tax evasion. *European Economic Review*, *33*, 797–805.

Gordon, R. H., & Cullen, J. B. (2012). Income redistribution in a federal system of governments. *Journal of Public Economics*, *96*, 1100–1109.

Görges, L., & Nosenzo, D. (2020). Social norms and the labor market. In K. F. Zimmermann (Ed.), *Handbook of labor, human resources and population economics* (pp. 1–26). Spring International Publishing.

Grodner, A., & Kniesner, T. J. (2006). Social interactions in labor supply. *Journal of the European Economic Association*, *4*, 1226–1248.

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

Hellwig, M. F. (2009). A note on Deaton's theorem on the undesirability of nonuniform excise taxation. *Economics Letters*, *105*, 186–188.

Hellwig, M. F. (2010). A generalization of the Atkinson-Stiglitz (1976) theorem on the undesirability of nonuniform excise taxation. *Economics Letters*, *108*, 156–158.

Hochman, H. M., & Rodgers, J. D. (1969). Pareto optimal redistribution. *American Economic Review*, *59*, 542–557.

Hosios, A. J. (1990). On the efficiency of matching and related models of search and unemployment. *Review of Economic Studies*, *57*, 279–298.

Hungerbühler, M., & Lehmann, E. (2009). On the optimality of a minimum wage: New insights from optimal tax theory. *Journal of Public Economics*, *93*, 464–481.

Hungerbühler, M., Lehmann, E., Parmentier, A., & Van der Linden, B. (2006). Optimal redistributive taxation in a search equilibrium model. *Review of Economic Studies*, *73*, 743–767.

Immonen, R., Kanbur, R., Keen, M., & Tuomala, M. (1998). Tagging and taxing: The optimal use of categorical and income information in designing tax/transfer schemes. *Economica*, *65*, 179–192.

Jacobs, B., & Lans Bovenberg, A. (2010). Human capital and optimal positive taxation of capital income. *International Tax and Public Finance*, *17*, 451–478.

Jacobs, B., & Lans Bovenberg, A. (2011). Optimal taxation of human capital and the earnings function. *Journal of Public Economic Theory*, *13*, 957–971.

Jacobs, B., Jongen, E. L. W., & Zoutman, F. T. (2017). Revealed social preferences of Dutch political parties. *Journal of Public Economics*, *156*, 81–100.

Jacquet, L., & Van der Linden, B. (2006). The normative analysis of tagging revisited: Dealing with stigmatization. *FinanzArchiv*, *62*, 168–198.

Jacquet, L., Lehmann, E., & Van der Linden, B. (2013). The optimal marginal tax rates with both extensive and intensive responses. *Journal of Economic Theory*, *148*, 1770–1805.

Jacquet, L., Lehmann, E., & Van der Linden, B. (2014). Optimal income taxation with Kalai wage bargaining and endogenous participation. *Social Choice and Welfare*, *42*, 381–402.

Kahneman, D., Knetsch, J. L., & Thaler, R. (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review*, *76*, 728–741.

Kanbur, R., & Tuomala, M. (1994). Inherent inequality and the optimal graduation of marginal tax rates. *The Scandinavian Journal of Economics*, *96*, 275–282.

Kanbur, R., & Tuomala, M. (2013). Relativity, inequality, and optimal nonlinear income taxation. *International Economic Review*, *54*, 1199–1217.

Kaplow, L. (2001). Horizontal equity: New measures, unclear principles. In K. Hassett & G. Hubbard (Eds.), *Inequality and tax policy* (pp. 75–97). American Enterprise Institute.

Kaplow, L. (2008). *The theory of taxation and public economics*. Princeton University Press.

Karni, E. (1999). Optimal unemployment insurance: A survey. *Southern Economic Journal*, *66*, 442–465.

Kasy, M. (2018). *Why a universal basic income is better than subsidies of low-wage work* (Data for Progress Working Paper). `http://filesforprogress.org/pdfs/UBI_EITC_Kasy_DFP_Working_Paper.pdf`

Keane, M. P. (2011). Labor supply and taxes: A survey. *Journal of Economic Literature*, *49*, 961–1075.

Kesselman, J., & Garfinkel, I. (1978). Professor Friedman, meet Lady Rhys-Williams: NIT vs CIT. *Journal of Public Economics*, *10*, 179–216.

King, M. L., Jr. (1967). *Where do we go from here: Chaos or community?* Harper and Row.

Kleven, H. J. (2014). How can Scandinavians tax so much? *Journal of Economic Perspectives*, *28*, 77–98.

Kleven, H. J. (2019). *The EITC and the extensive margin: A reappraisal* (NBER Working Paper

No. 26405).

Kleven, H. J., & Kopczuk, W. (2011). Transfer program complexity and the take up of social benefits. *American Economic Journal: Economic Policy*, *3*, 54–90.

Kleven, H. J., & Schultz, E. A. (2014). Estimating taxable income responses using Danish tax reforms. *American Economic Journal: Economic Policy*, *6*, 271–301.

Koebel, K., & Pohler, D. (2019). Expanding the Canada Workers Benefit to design a guaranteed basic income. *Canadian Public Policy*, *45*, 283–309.

Krause, A. (2006). Redistributive taxation and public education. *Journal of Public Economic Theory*, *8*, 807–819.

Krause, A. (2009). Optimal nonlinear income taxation with learning-by-doing. *Journal of Public Economics*, *93*, 1098–1110.

Kroft, K., Kucko, K., Lehmann, E., & Schmieder, J. (2020). Optimal income taxation with unemployment and wage responses: A sufficient statistics approach. *American Economic Journal: Economic Policy 12*(1), 254–292.

Kuziemko, I., Norton, M. I., Saez, E., & Stantcheva, S. (2015). How elastic are preferences for redistribution? Evidence from randomized survey experiments. *American Economic Review*, *105*, 1478–1508.

Landais, C., Michaillat, P., & Saez, E. (2018). A macroeconomic approach to optimal unemployment insurance: Theory. *American Economic Journal: Economic Policy*, *10*, 152–181.

Lans Bovenberg, A., & Jacobs, B. (2005). Redistribution and education subsidies are Siamese twins. *Journal of Public Economics*, *89*, 2005–2035.

Lavecchia, A. (2020). Minimum wage policy with optimal taxes and unemployment. *Journal of Public Economics*, *190*, Article 104228.

Layard, R. (1980). Human satisfactions and public policy. *The Economic Journal*, *90*, 737–750.

Lee, D., & Saez, E. (2012). Optimal minimum wage policy in competitive labor markets. *Journal of Public Economics*, *96*, 739–749.

Lehmann, E., Parmentier, A., & Van der Linden, B. (2011). Optimal income taxation with endogenous participation and search unemployment. *Journal of Public Economics*, *95*, 1523–1537.

Lockwood, B. B. (2020). Optimal income taxation with present bias. *American Economic Journal: Economic Policy*, *12*(4), 298–327.

Lockwood, B. B., & Weinzierl, M. (2016). Positive and normative judgments implicit in U.S. tax policy, and the costs of unequal growth and recessions. *Journal of Monetary Economics*, *77*, 30–47.

Luttmer, E. F. P., & Singhal, M. (2014). Tax morale. *Journal of Economic Perspectives*, *28*(4), 149–168.

Mankiw, N. G., & Weinzierl, M. (2010). The optimal taxation of height: A case study of utilitarian income redistribution. *American Economic Journal: Economic Policy*, *2*, 155–176.

Marceau, N., & Boadway, R. (1994). Minimum wage legislation and unemployment insurance as instruments for redistribution. *Scandinavian Journal of Economics*, *96*(1), 67–81.

Meade, J. (1936). *An introduction to economic analysis and policy*. Oxford University Press.

Meltzer, A. H., & Richard, S. F. (1981). A rational theory of the size of government. *Journal of Political Economy*, *89*, 914–927.

Mill, J. S. (1987). *Principles of Political Economy* (2nd ed.). Augustus Kelley. (Original work published 1849)

Milligan, K., & Smart, M. (2015). Taxation and top incomes in Canada. *Canadian Journal of Economics. Revue Canadienne d'Economique*, *48*, 655–681.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *Review of Economic Studies*, *38*, 175–208.

Mirrlees, J. A. (2007, July 12–13). *Taxation of gifts and bequests* [Unpublished slides]. James Meade Centenary Conference, London, UK.

More, T. (1963). *Utopia* (P. Turner, Trans.). Penguin. (Original work published 1516)

Myles, G. (1995). *Public economics*. Cambridge University Press.

Nishimura, Y. (20003). Optimal commodity taxation for the reduction of envy. *Social Choice and Welfare*, *21*(3), 501–527.

Office of the Parliamentary Budget Officer. (2016). The fiscal and distributional impact of changes to the federal personal income tax regime. `https://www.pbo-dpb.gc.ca/web/default/files/Documents/Reports/2016/PIT/PIT_EN.pdf`

Olivera, J. (2015). Preferences for redistribution in Europe. *IZA Journal of European Labor Studies*, *4*, Article 14.

Osberg, L. (2018). *The age of increasing inequality*. James Lorimer and Company.

Osborne, M. J., & Slivinski, A. (1996). A model of political competition with citizen-candidates. *Quarterly Journal of Economics*, *111*, 65–96.

Oswald, A. (1983). Altruism, jealousy and the theory of optimal non-linear taxation. *Journal of*

Public Economics, *20*, 77–87.

Paine, T. (1796). *Agrarian justice*. Evans and Bone.

Parsons, D. O. (1996). Imperfect "tagging" in social insurance programs. *Journal of Public Economics*, *62*, 183–207.

Piketty, T., & Saez, E. (2013). Optimal labor income taxation. In A. Auerbach, R. Chetty, M. Feldstein, & E. Saez (Eds.), *Handbook of public economics* (Vol. 5, pp. 391–474). North-Holland Publishing Company.

Pissarides, C. A. (2000). *Equilibrium unemployment theory* (2nd ed.). MIT Press.

Prendergast, C. (2007). The motivation and bias of bureaucrats. *American Economic Review*, *97*, 180–196.

Reich, R. (2016, September 29). Why we'll need a universal basic income. *Robert Reich*. `http://robertreich.org/post/151111696805`

Roemer, J. E. (1998). *Equality of opportunity*. Harvard University Press.

Rogerson, R., Shimer, R., & Wright, R. (2005). Search-theoretic models of the labor market: A survey. *Journal of Economic Literature*, *43*, 959–988.

Rothstein, J. (2010). Is the EITC as good as an NIT? Conditional cash transfers and tax incidence. *American Economic Policy: Economic Policy*, *2*, 177–208.

Rothschild, C., & Scheuer, F. (2013). Redistributive taxation in the Roy model. *Quarterly Journal of Economics*, *128*, 623–668.

Russell, B. (1918). *Roads to freedom: Socialism, anarchism and syndicalism*. Unwin Books.

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

Saez, E. (2002). Optimal income transfer programs: Intensive vs extensive labor supply responses. *Quarterly Journal of Economics*, *117*, 1039–1073.

Saez, E. (2010). Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy*, *2*, 180–212.

Saez, E., & Stantcheva, S. (2016). Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, *106*, 24–45.

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

Saez, E., & Veall, M. (2005). The evolution of high incomes in Northern America: Lessons from Canadian evidence. *American Economic Review*, *95*, 831–849.

Segal, H. D. (2016). *Finding a better way: A basic income pilot project for Ontario* [Discussion paper]. Government of Ontario. `https://www.ontario.ca/page/finding-better-way-basic-income-pilot-project-ontario`

Sen, A. (1973). *On economic inequality*. Clarendon Press.

Shapiro, C., & Stiglitz, J. E. (1984). Equilibrium unemployment as a worker discipline device. *American Economic Review*, *74*, 433–444.

Sheshinski, E. (1972). The optimal linear income tax. *Review of Economic Studies*, *39*(3), 297–302.

Sheshinski, E. (1989). Note on the shape of the optimum income tax schedule. *Journal of Public Economics*, *40*, 201–215.

Sillamaa, M.-A., & Veall, M. (2001). The effect of marginal tax rates on taxable income: A panel study of the 1988 tax flattening in Canada. *Journal of Public Economics*, *80*, 341–356.

Simon, H. A. (2000). UBI and the flat tax. In P. Van Parisj (Ed.), *What's wrong with a free lunch?* (pp. 34–38). Beacon Press.

Slemrod, J., Yitzhaki, S., Mayshar, J., & Lundholm, M. (1994). The optimal two-bracket linear income tax. *Journal of Public Economics*, *53*, 269–290.

Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association*, *13*, 130–167.

Standing, G. (2017). *Basic income: A guide for the open-minded*. Yale University Press.

Stantcheva, S. (2015). *Learning and (or) doing: Human capital investments and optimal taxation* (NBER Working Paper No. 21381).

Strawczynski, M. (1998). Social insurance and the optimum piecewise linear income tax. *Journal of Public Economics*, *69*(3), 371–388.

Stevens, H., & Simpson, W. (2017). Toward a national universal guaranteed basic income. *Canadian Public Policy*, *43*, 120–139.

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

Tanninen, H., Tuomala, M., & Tuominen, E. (2019). *Inequality and optimal redistribution*. Cambridge University Press.

Tarroux, B. (2019). The value of tax progressivity: Evidence from survey experiments. *Journal of Public Economics*, *179*, Article 104068.

Tinbergen, J. (1975). *Income distribution*. North-Holland Publishing Company.

Tobin, J. (1966). The case for an income guarantee. *The Public Interest*, *Summer*, 31–41.

Topel, R., & Welch, F. (1980). Unemployment insurance: Survey and extensions. *Economica*, *47*, 351–379.

Tuomala, M. (1990). *Optimal income tax and redistribution*. Clarendon Press.

Tuomala, M. (2016). *Optimal redistributive taxation*. Oxford University Press.

Van Parijs, P. (1995). *Real freedom for all*. Oxford University Press.

Veall, M. (2012). Top income shares in Canada: Recent trends and policy implications. *Canadian Journal of Economics. Revue Canadienne d'Economique*, *45*, 1247–1272.

von Hayek, F. A. (1979). *Law, legislation and liberty: Vol. 3. The political order of a free people*. University of Chicago Press.

Weinberg, B. A., Reagan, P. B., & Yankow, J. J. (2004). Do neighborhoods affect hours worked? Evidence from longitudinal data, *Journal of Labor Economics*, *22*, 891–924.

Weinzierl, M. C. (2011). The surprising power of age-dependent taxes. *Review of Economic Studies*, *78*, 1490–1518.

Weinzierl, M. (2014). The promise of positive optimal taxation: Normative diversity and a role for equal sacrifice. *Journal of Public Economics*, *118*, 128–142.

Wenzel, M. (2005). Motivation or rationalisation? Causal relations between ethics, norms and tax compliance. *Journal of Economic Psychology*, *26*, 491–508.

Woittiez, I., & Kapteyn, A. (1998). Social interactions and habit formation in a model of female labour supply. *Journal of Public Economics*, *70*, 185–205.

Figure 1

Linear Progressive Income Tax

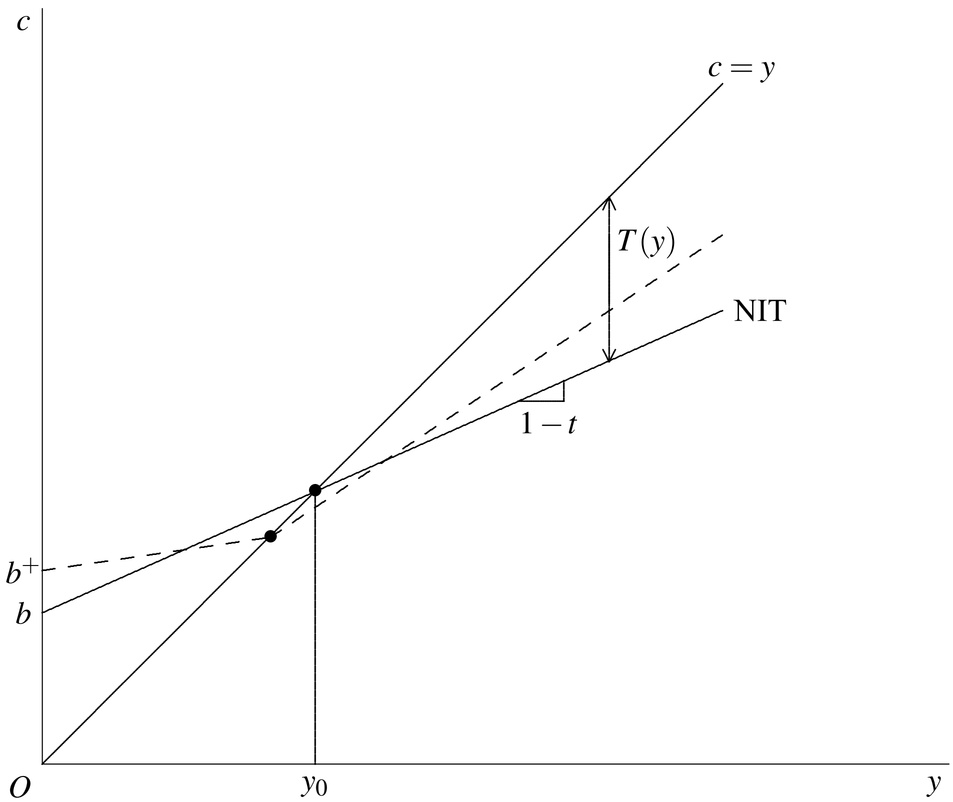

Figure 2

Piecewise Linear Progressive Tax

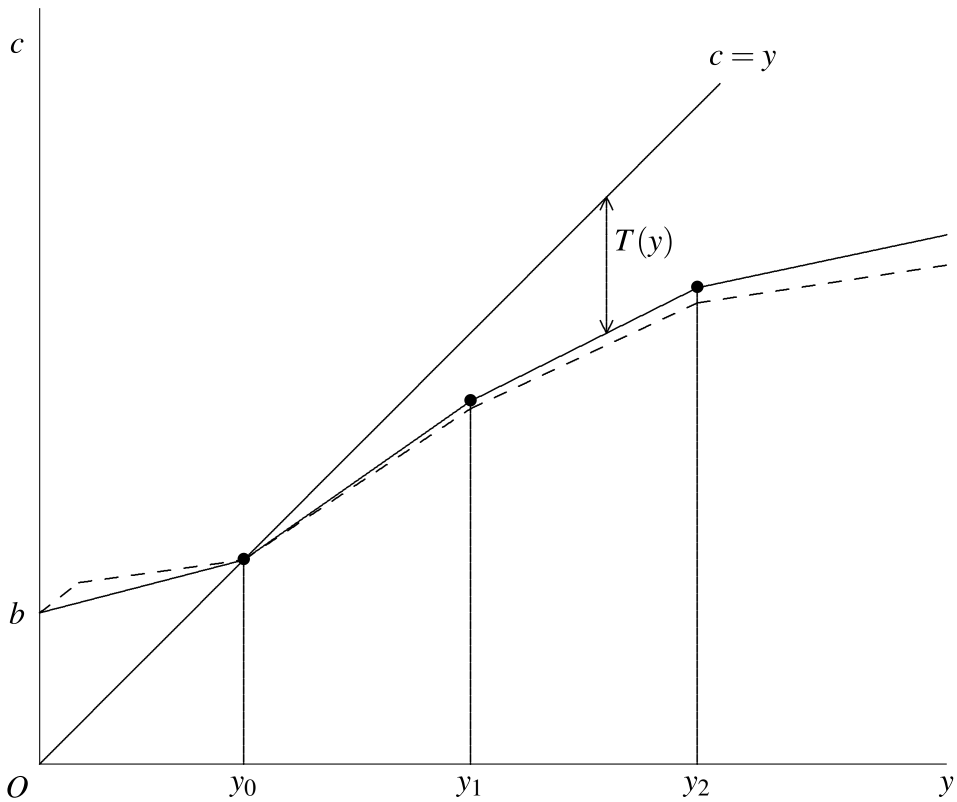

Figure 3

Nonlinear Income Tax With Intensive Labour Supply

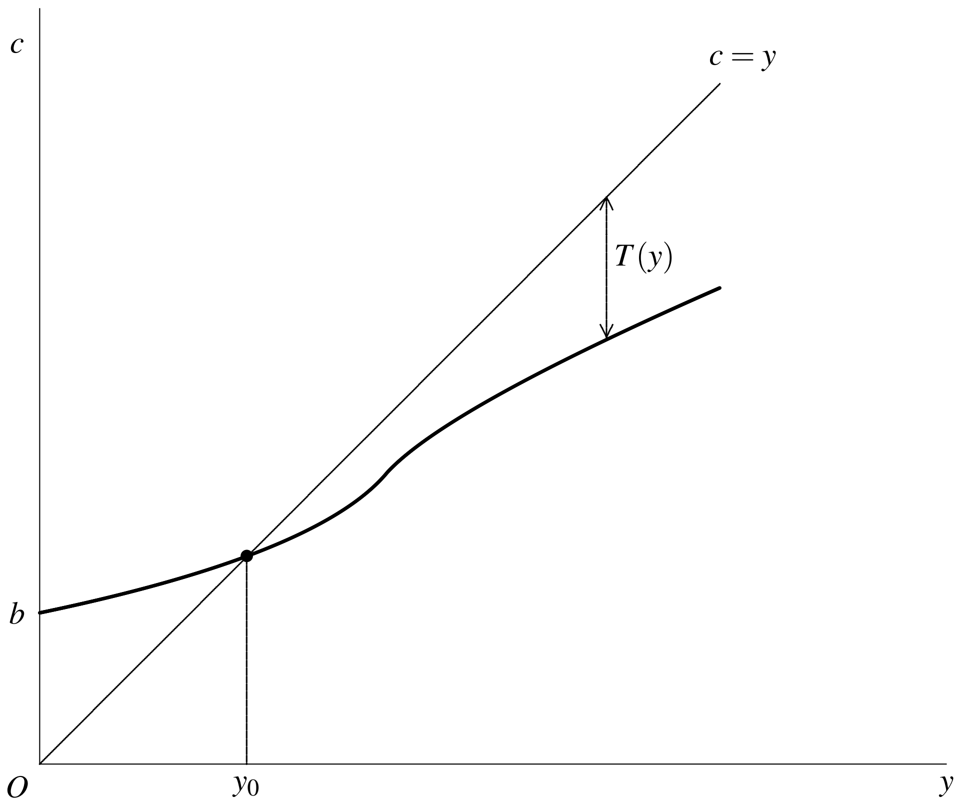

Figure 4

Nonlinear Income Tax With Extensive Labour Supply

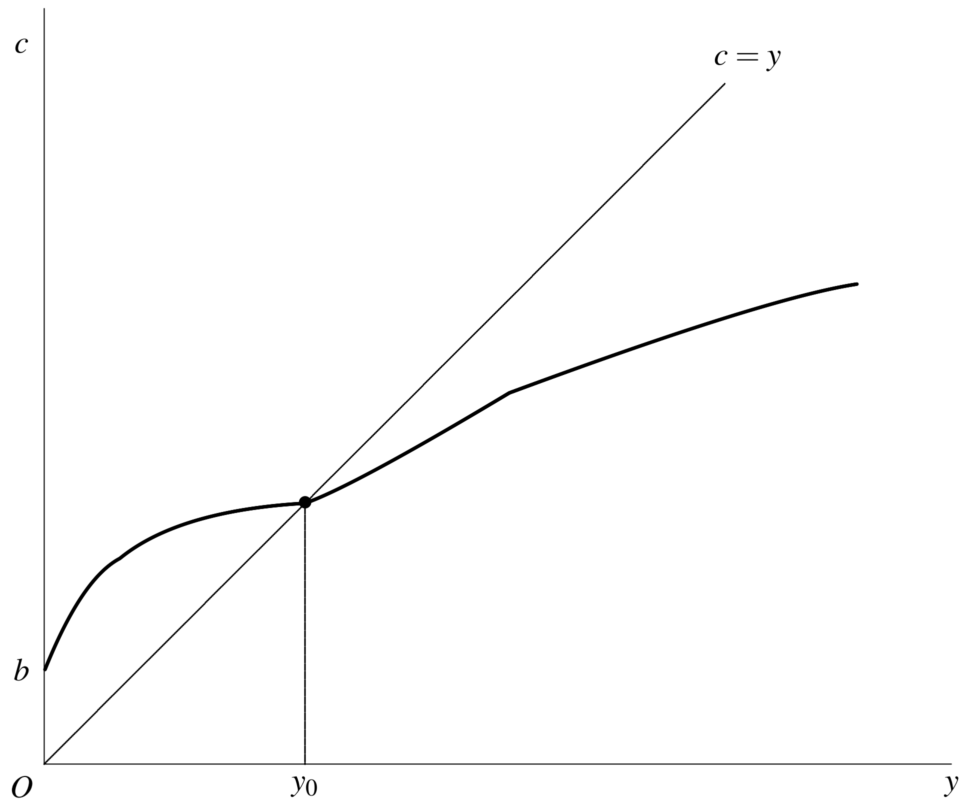

Figure 5

Champernowne vs. Lognormal

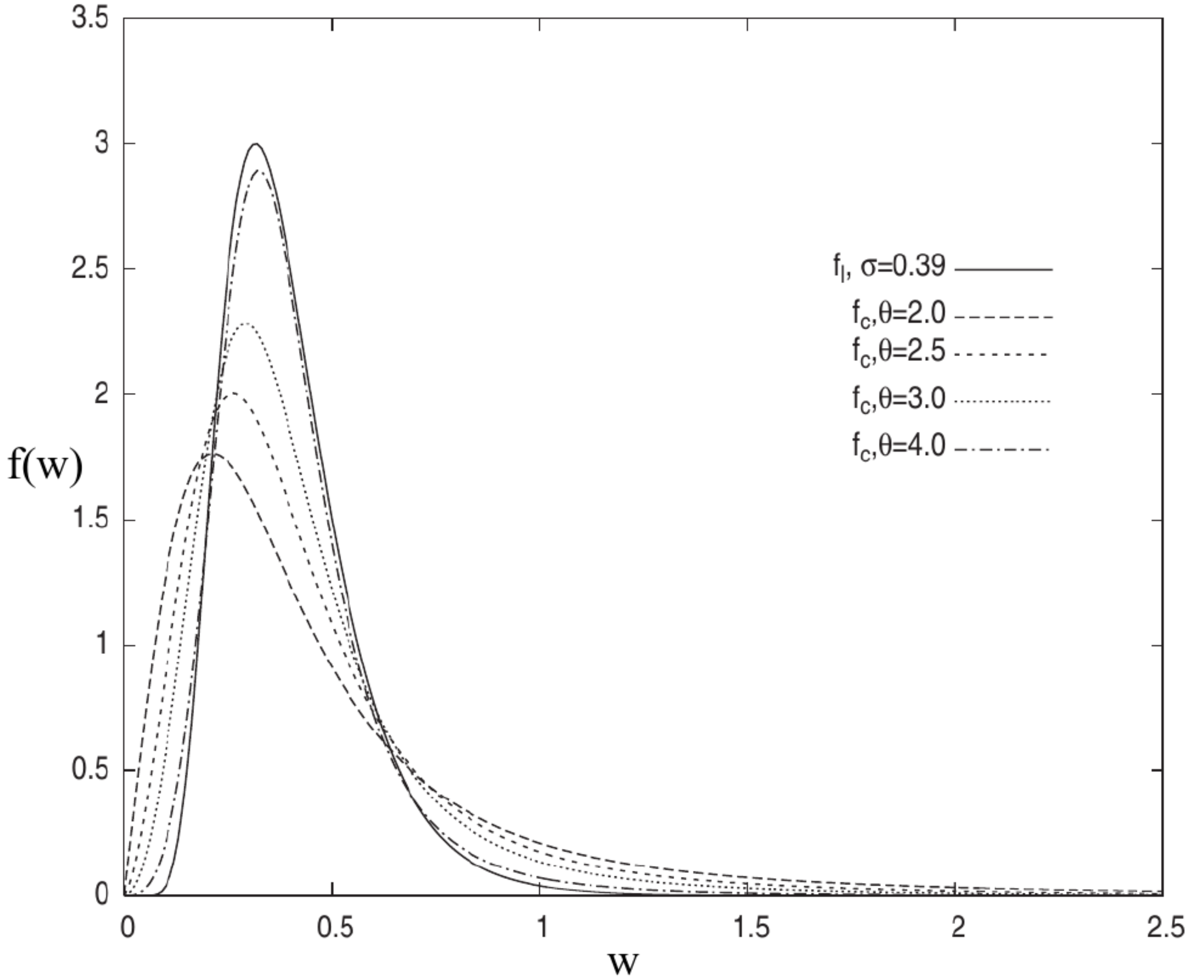

Figure 6

Tuomala, Average Tax Rates

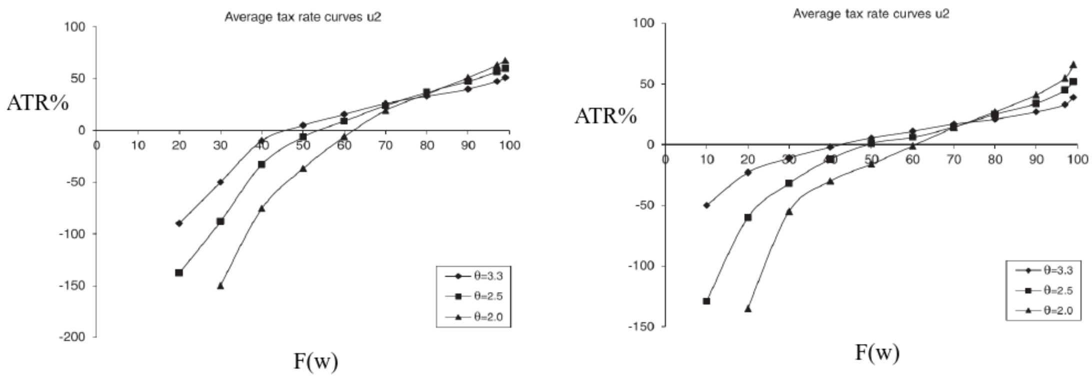

(a) Maximin (b) Utilitarian

Figure 7

Tuomala, Marginal Tax Rates (Champernowne)

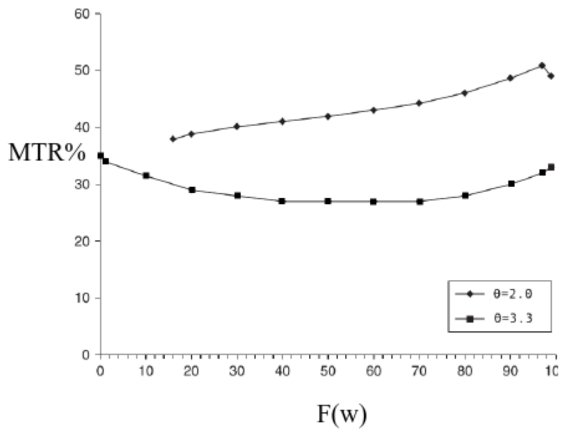

(a) Effect of Pretax inequality

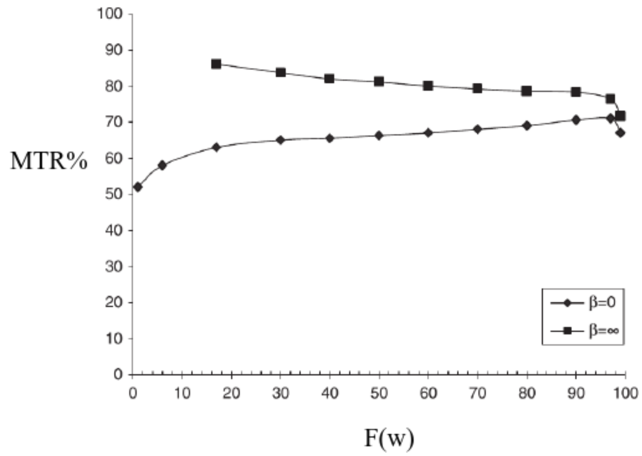

(b) Utilitarian vs Maximin, u2