

Income Taxation with Frictional Labor Supply

Nicolas Werquin*

November 16, 2014

Abstract

This paper studies the effects of taxes on labor income and social welfare in an environment where labor supply choices are constrained by adjustment frictions. I analyze a dynamic model in which individuals choose their labor supply on the intensive margin as a function of their stochastic idiosyncratic productivity shocks and the non-linear tax schedule. Agents incur a fixed cost of adjusting their labor supply in response to productivity or tax changes, which can be thought of as the cost of searching for a new job. In the frictionless economy, I derive sufficient statistic formulas for the long-run effects of local tax reforms on social welfare. In the frictional model, the first main result is that, for a given labor income elasticity, the long-run effects of tax changes on social welfare differ significantly from those in the frictionless economy. The frictionless model ignores the heterogeneity in the utility of individuals who earn the same income level, and thus systematically underestimates the welfare costs of raising marginal tax rates. Moreover, this distribution is endogenous to taxes, leading to higher welfare gains from a budget-neutral increase in the progressivity of the tax schedule. The second main result is that the three-year elasticity of labor income to marginal tax rates typically estimated in the data may underestimate the long-run aggregate elasticity when frictions are present, more so when the proportion of exogenously received to endogenously chosen labor supply adjustments increases.

*Yale University. Email: nicolas.werquin@yale.edu. Please check for updates at <http://nwerquin.com>. I am extremely grateful to Mikhail Golosov, Giuseppe Moscarini, and Aleh Tsyvinski for their invaluable guidance and support, and to Joe Altonji, Costas Arkolakis, Eduardo Engel, Tibor Heumann, Johannes Hörner, Costas Meghir, Michael Peters, Tommaso Porzio, and Emmanuel Saez for very helpful discussions at the various stages of this project. I also thank Shameel Ahmad, Antoine Arnoud, David Childers, Gabriele Foà, Sebastian Heise, Dominik Sachs, Ali Shourideh, Stefanie Stantcheva, Yujung Whang, and the participants at the Yale macroeconomics and public economics seminars for helpful comments and suggestions.

1 Introduction

A large body of empirical evidence shows that the adjustment of labor supply in response to productivity or tax changes is subject to frictions. Yet despite this large and growing empirical literature, there is little theoretical work that explicitly incorporates such adjustment frictions into models of income taxation. What are the consequences of adjustment frictions for the social welfare effects of raising marginal tax rates or the degree of progressivity of the tax schedule? Second, after a tax change, how fast does the economy as a whole adjust, and what are the consequences for the long-run elasticity of labor income to marginal tax rates?

This paper studies a model in which individuals face a fixed cost of adjusting their hours of work or taxable income in response to productivity or tax changes. First, I set up a dynamic frictionless model of the economy that is tractable enough to be solved entirely in closed form, and characterize the long-run effects of taxes in this environment. Second, I add adjustment frictions to this model and analyze the long-run effects of tax changes on social welfare. I show that in the presence of frictions the welfare costs of raising marginal tax rates and the degree of progressivity differ significantly from those in the frictionless model. Third, I analyze the impulse response function of aggregate income to a tax change until convergence to the steady state and show that the three-year elasticity of labor income to marginal tax rates typically estimated in the data may underestimate the long-run effects of tax changes.

Specifically, I analyze a dynamic continuous-time model in which individuals choose their labor supply on the intensive margin as a function of their stochastic idiosyncratic productivity shocks and the non-linear tax schedule. The tax instruments that are available to the government are uniform changes in the marginal tax rates, and changes in the degree of progressivity of the tax schedule. I first study the steady state of the model without frictions. The results I derive in this environment are interesting on their own as I provide a full closed-form characterization of the individual behavior and the aggregate income distributions as a function of the parameters of the individual shock processes and of the tax schedule. Individual responses to taxes in this model are similar to those in the canonical *static* model of income taxation, which makes my results easily comparable to this literature. However, the productivity and income distributions that I obtain are not exogenously given as in the standard model, but endogenously determined as the stationary outcome of a dynamic process given the tax function. Moreover, the individual discount rate that is relevant to compute the present discounted value of utility depends on the growth rate of future consumption and hence is endogenous to taxes. This effect of taxes on social welfare through the growth rate of future consumption is not captured in the static taxation model. This allows me to interpret in a clean way this economy as a benchmark framework to analyze the question of the *long-run* effects of taxes in a frictionless environment. In this model, I derive “sufficient statistic” formulas (see, e.g., Saez 2001 and Golosov, Tsyvinski, and Werquin 2014) that characterize the effects of local tax reforms on social welfare. The key parameters that determine the welfare implications of raising taxes are: (*i*) the *elasticity of labor income* to

marginal tax rates, which measures the individual behavioral responses to higher taxes, and (ii) the *marginal social welfare weights*, which summarize the redistributive tastes of the government and the welfare gains and losses of changing taxes.

Next, I analyze the *frictional* model where agents also incur a fixed cost of adjusting their labor supply or taxable income in response to productivity or tax changes. This fixed cost can be thought of for example as the cost of searching for a new job. In addition, individuals randomly receive exogenous adjustment opportunities at no cost. The case where there are no such opportunities and a positive fixed cost represents an environment where all adjustments are endogenous and driven by labor supply variables, namely, the evolution of productivity and taxes. Intuitively, individuals would like to work and earn more in response to an increase in their productivity, but to do so they have to pay the search cost and find a new job. In this framework, individuals are inactive most of the time and choose to adjust whenever their optimal income (which they would choose in the absence of frictions) gets far enough from their current actual income. This model is similar to the (S, s) (or (L, c, U)) models analyzed in the operations research, monetary and investment literatures (see, e.g., Harrison, Sellke, and Taylor 1985, Dixit and Pindyck 1994, Stokey 2008). At the other extreme, if the fixed cost is infinite, individuals search permanently but can only adjust upon receiving exogenous costless opportunities to do so (job offers); this represents the case where all adjustments are driven by labor demand. This model is similar to that of Calvo (1983). In the general framework, which nests the continuum of intermediate cases between these two polar setups, both margins are active, and the relative size of the arrival rate of exogenous opportunities and the fixed adjustment cost determines the relative strength of the labor demand side of the market in constraining individual income decisions. This framework is similar to the CalvoPlus setup that Nakamura and Steinsson (2010) and Alvarez, Le Bihan, and Lippi (2014) studied in the monetary literature. This paper brings these tools to the theory of taxation. An important difference is that unlike the models of money, taxes have real effects in the long run; those effects are the subject of this paper.

I fully characterize the individual and aggregate behavior of the frictional economy. Individuals have two decision variables. The first is their *desired* or frictionless income, which follows the same dynamics as in the frictionless model. The second is the *deviation* of their actual income away from their desired income. I show that this variable follows dynamics given by an (L, c, U) policy, i.e., adjustment occurs when the deviation crosses an upper threshold (the individual works “too much”) or a lower threshold (he works “too little”), unless a costless adjustment opportunity is received before the deviation reaches the boundaries of the inaction region. Importantly, the optimal size of the inaction region and the frequency of adjustment are endogenous to tax policy. I characterize the distributions of incomes and deviations that are induced by the behavior of individuals.

I then analyze the effects of taxes on *long-run* social welfare in the frictional model, and compare these effects to those in the frictionless model. In a model where the only government’s tax instrument is to raise or lower the marginal tax rates uniformly, I show that the effects

of taxes on long-run social welfare are given by a sufficient statistic formula which is formally identical to that in the frictionless model. In particular, the long-run elasticity that is relevant for this problem is the individual *structural* (frictionless) labor income elasticity. Intuitively, in the long run all individuals have had the time to fully adjust to the new tax system, and the behavior of aggregate income is driven by same the structural elasticity parameter as in the frictionless model. However, the relevant social welfare weights which characterize the redistributive tastes of the government are different from those in the frictionless model, even in the long run. The presence of adjustment frictions implies that at any point in time, individuals who earn the same income differ in their utility, as the least productive of them must provide higher effort than the others to earn the same income. Hence there is a non-degenerate distribution of deviations at every income level. By ignoring this heterogeneity and treating the population earning the same income as a representative agent, the frictionless model systematically underestimates the costs of raising marginal tax rates. The key friction is that the labor income taxes cannot disentangle individuals who differ only in their unobservable effort or productivity. If the government is redistributive enough, i.e. if the social welfare function is concave enough, this implies that the marginal social welfare weights are higher at every level of income, and therefore that the welfare costs of raising the marginal tax rates are higher than in the frictionless model. I further quantify the size of this effect, and show that the costs of raising taxes can be up to twice as large in the frictional environment for reasonable social welfare functions.

Next, I study the effects of changing the progressivity of the tax schedule in order to redistribute income across individuals, subject to the government budget constraint. I show that the sufficient statistic formula that characterizes the welfare effects of those tax reforms in the frictionless model no longer holds in the presence of frictions, even replacing the welfare weights with the relevant frictional ones. First, the tax rates affect not only the optimal desired labor supply, but also the frequency of adjustment. When an increase in progressivity leads to a slower job turnover, the dispersion of individuals around their optimum increases, which tends to reduce social welfare. This deformation of the income distribution also implies that the structural labor income elasticity is not a sufficient statistic for the effects of taxes, so that the population earning a given income level cannot be characterized by a representative agent, even as far as the effects on tax revenue are concerned. Second, the distribution of heterogeneous utilities within an income group is itself endogenous to tax policy. I show that this endogeneity implies higher welfare gains of increasing the rate of progressivity than in the frictionless model, if the planner is redistributive enough. Numerically these effects are modest, yet non-negligible. This result shows that simply acknowledging the fact that the population is heterogeneous at each income level without microfounding this distribution is insufficient to calculate the effects of taxes: a structural model is needed to fully account for their impact on social welfare.

Finally, I analyze the impulse response functions of aggregate income following a tax change until convergence to the steady state. The micro literature (see, e.g., Saez, Slemrod, and Giertz 2012 for a survey) typically estimates labor income elasticities at the horizon of three years

after a tax reform. In the presence of adjustment frictions, however, an important question is whether the long-run (say, ten-year) elasticity is substantially larger than the three-year elasticity. I provide bounds for the speed of adjustment of the economy in a calibrated version of the model. For a given frequency of individual income adjustments, the convergence of the aggregate economy is the fastest when there are no exogenous costless opportunities to adjust, for then the size of the inaction region is the narrowest and the individuals who adjust first are those whose labor income is the farthest from their optimum. This is a selection effect similar to that highlighted by Golosov and Lucas (2007). In this case, the three-year elasticity provides a close estimate to the long-run elasticity. On the other hand, the adjustment is the slowest when all adjustment opportunities are exogenously received rather than endogenously chosen, as in the Calvo (1983) model. In this case, the long-run elasticity is larger by a third relative to the three-year elasticity. Therefore, in the nested CalvoPlus model, when the proportion of exogenously received to endogenously chosen labor supply adjustments increases, i.e., the constraints imposed by the labor demand side of the market become more stringent relative to the purely labor supply driven decisions, the long-run aggregate labor income elasticity increasingly diverges from the three-year estimated elasticity.

Related literature. This paper is related to the empirical literature that points to the presence of frictions in the adjustment of labor supply. Altonji and Paxson (1992) show that changes in labor supply preferences have a much larger effect on hours of work when individual change jobs, suggesting that adjusting behavior entails substantial fixed costs. Other papers have argued that labor supply is constrained by adjustment costs and hours constraints, e.g., Cogan (1981), Altonji and Paxson (1988), Dickens and Lundberg (1993), Chetty, Friedman, Olsen, and Pistaferri (2011), Gelber, Jones, and Sacks (2013). My primary contribution is to model explicitly these fixed costs into a dynamic model of income taxation and derive the consequences for the welfare effects of taxes.

This paper also relates to the theoretical literature on adjustment frictions in a taxation context. Chetty, Looney, and Kroft (2009) propose a model of bounded rationality where individuals' responses to taxes are affected by tax salience, and show that this feature affects the calculation of the impact of taxes on social welfare, an insight related to the effect found in this paper. Chetty, Friedman, Olsen, and Pistaferri (2011) study a model in which labor supply is subject to search costs and jobs are characterized by hours constraints. Saez (2002), Choné and Laroque (2011), Jacquet, Lehmann, and Van der Linden (2013), Shourideh and Troshkin (2014) study optimal taxation problems where labor supply is set on the extensive margin, i.e., individuals face a fixed cost of working. These models are primarily static, and cannot capture the dynamic decisions of individuals to adjust based on their option value of waiting. Modeling this dynamic behavior is important, as it allows us to endogenize the distributions of imbalances (deviations away from optimal choices) at each income level, on which the effects of taxes crucially depend. Moreover, there are many more sources of extensive margin decisions than

the choice of participation (e.g., retirement) in practice, and my model allows to analyze those related to the discrete choices of hours and incomes of employed individuals.

My frictionless model is related to that of Heathcote, Storesletten, and Violante (2014). They also restrict the set of available tax instruments to two-parameter schedules, and analyze the effects of progressivity on social welfare in a model with imperfect private insurance and investment in skills. My frictionless model is simpler and allows me to get extremely transparent closed-form expressions for the effects of taxes on the income distributions and social welfare. The income distributions I obtain endogenously have been argued to be a very good fit of the actual empirical distributions (Reed 2003, Reed and Jorgensen 2004, Toda 2012). I naturally obtain income distributions that have Pareto tails using a random growth process for individual productivities and incomes (see, e.g., Gabaix 2009). A large empirical literature estimates precisely this type of earnings processes (e.g., Meghir and Pistaferri, 2004 and 2011 for a survey).

Most importantly, this model allows me to introduce labor supply adjustment frictions in a tractable way. The frictional model is related to the literature on impulse control models. Dixit and Pindyck (2004) and Stokey (2008) summarize many applications of these methods to economics. Bertola and Caballero (1994), Grossman and Laroque (1990), Caballero and Engel (1999), and more recently Alvarez and Lippi (2013) and Alvarez, Le Bihan, and Lippi (2014), among many others, have made important theoretical contributions to this literature, on which this paper builds. I bring this literature to the field of public finance, and analyze the long run effects of taxes in this framework.

My paper is also related to the large literature on the labor income elasticities. The elasticity this paper is concerned about is the Hicksian (steady state) elasticity. The micro literature typically finds small elasticities (0.3 or lower, see Saez, Slemrod, and Giertz 2012 for a survey), while the macro literature finds elasticities closer to 1 (see Keane and Rogerson 2012 for a survey). Rogerson and Wallenius (2009, 2013) and Ljungqvist and Sargent (2011) argue that the small micro and large macro elasticities can be reconciled if the primary margin of adjustment of labor supply is the choice of career length (mostly retirement) rather than hours conditional on participation. Chetty, Guren, Manoli, and Weber (2012) criticize this view, and Chetty (2012) argue that adjustment frictions can explain the difference between the micro and macro Hicksian elasticities. Moreover, Holmlund and Söderström (2008) argue that the short-run and the long-run elasticities may differ. My paper endogenously generates extensive margin responses for employed individuals, and studies whether the long-run elasticities differ significantly from the short-run elasticities.

Finally, my analysis is related to the taxation literature. As in Mirrlees (1971) and Diamond (1998), I model labor supply choices on the intensive margin when individuals choose their income based on their exogenous productivity; the government's only available instrument is to tax labor income. I interpret the static taxation framework as a model of the long-run by embedding it into a dynamic environment, in which the income distribution is the endogenous stationary outcome of the underlying idiosyncratic shock process at the micro level. There

is a large literature that derives sufficient statistic formulas for the effects of taxation: e.g., Chetty (2009) for a general exposition, Saez (2001) in the static setting, Golosov, Tsyvinski, and Werquin (2014) in the dynamic setting. These formulas are valid for a very large class of models and underlying functional forms for the utility functions, the sources of heterogeneity, etc. However, these models generally assume that labor supply can always be set optimally at no cost. In this paper I revisit these sufficient statistic formulas and show that they may not hold in the presence of adjustment frictions.

The structure of the paper is as follows. I set up the environment and solve the frictionless model in Section 2. I analyze the long run effects of taxation in this model in Section 3. I introduce and analyze the frictional model in Section 4. I study the long run and short run effects of taxes in this environment in Section 5. Section 6 concludes. The proofs of the results are gathered in the Appendix.

2 A long-run income taxation model

In this section I set up a model of taxation and frictionless labor supply. I derive the individuals' optimal behavior and the resulting income distributions in closed form as a function of the parameters of the tax schedule. This framework can be naturally interpreted as a *long-run* model of taxation, in which the steady-state income distribution is the endogenous stationary outcome of an underlying dynamic process.

2.1 Environment

Individuals. There is a continuum of mass one of individuals in the economy. Time is continuous. The idiosyncratic productivity θ_t of an individual is exogenous and evolves stochastically according to a geometric Brownian motion, i.e., a random growth process, with expected growth rate $g_\theta = \mu_\theta + \sigma_\theta^2/2$ and volatility σ_θ . That is,

$$d \ln \theta_t = \mu_\theta dt + \sigma_\theta d\mathcal{W}_t, \tag{1}$$

where \mathcal{W}_t is a Wiener process. Let $\{\mathcal{F}_t\}$ denote the filtration generated by \mathcal{W}_t . This specification implies that the log-productivity process has a unit root, i.e., productivity shocks are permanent. Individuals observe their exogenous productivity θ_t at every instant t , and choose their (endogenous) labor supply given θ_t .

An individual with productivity θ who provides e units of effort has an effective labor supply y equal to

$$y = \theta \times e.$$

I abstract from general equilibrium considerations and assume that the wage per efficiency

unit of labor in the economy is constant, $w = 1$, so that y represents interchangeably the individual's *hours of work* or his *taxable labor income*.¹ The utility function U is quasilinear in consumption c and the disutility of effort e is convex, specifically isoelastic with constant labor income elasticity $\varepsilon > 0$. Quasilinearity implies that without loss of generality consumption is equal at every instant to the disposable income, $c = y - T(y)$, where $T(\cdot)$ is the tax on labor income levied by the government.

The individual's endogenous choice variable is his effort e , or equivalently his taxable income $y \in \mathbb{R}_+$. The flow utility of an individual with productivity θ who earns income y can thus be written as

$$U(\theta, y) = y - T(y) - \frac{1}{1 + 1/\varepsilon} \left(\frac{y}{\theta}\right)^{1+1/\varepsilon}. \quad (2)$$

The labor force participation decision and the search activity of non-employed individuals are unaffected by tax policy. There is an exogenous and constant Poisson rate β of job creation (or "birth") and job destruction (or "death"), independent of an individual's productivity or income. An individual who re-enters the labor force draws a new productivity level θ_0 from an exogenous log-normal distribution with mean m_θ and variance s_θ^2 , i.e., $f_{\theta_0}(\cdot) \sim \log \mathcal{N}(m_\theta, s_\theta^2)$, and chooses his income $y_0(\theta_0)$ optimally.

Tax system. The government chooses a tax-and-transfer system $T(y)$. The tax system is restricted within a class of two-parameter schedules (see, e.g., Benabou, 2002; Heathcote, Storesletten, and Violante, 2014), defined as

$$T(y) = y - \frac{1 - \tau}{1 - p} y^{1-p}, \quad (3)$$

with $(\tau, p) \in \mathbb{R}_+ \times \mathbb{R}$. I denote the tax system interchangeably by T or $\{\tau, p\}$. The parameter p is the coefficient of marginal rate progression (see Musgrave and Thin, 1948). It is equal to the elasticity of the net-of-tax rate with respect to taxable income,

$$p = -\frac{d \ln(1 - T'(y))}{d \ln y}. \quad (4)$$

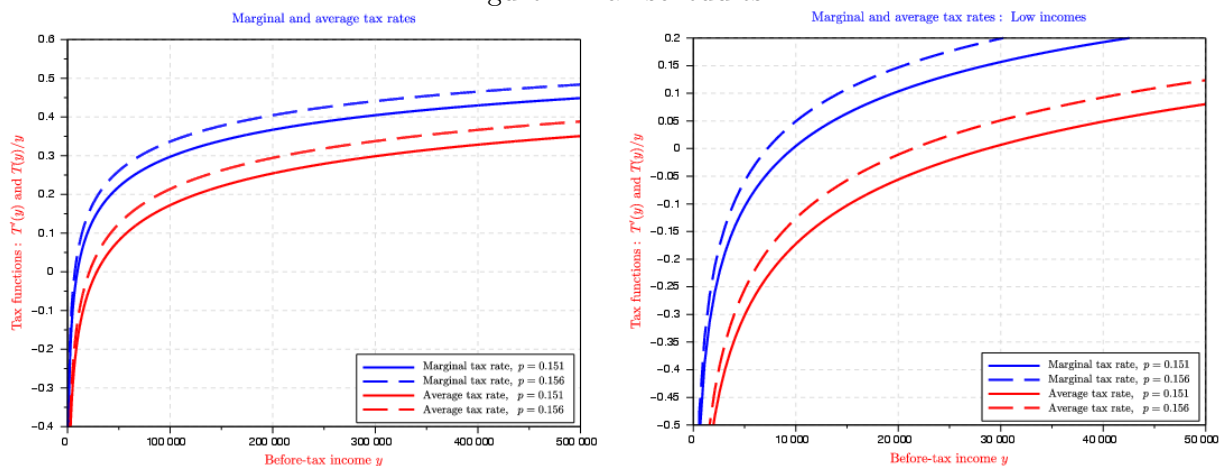
If $p = 0$, the income tax schedule is linear with constant marginal tax rate τ . If $p \in (0, 1)$, the ratio of the marginal tax rate to the average tax rate is $T'(y) / \{T(y)/y\} > 1$, so that the tax schedule is progressive. If $p < 0$, the tax schedule is regressive. Note that the marginal and the average tax rates are monotone in earnings, and that average tax rates are negative for incomes y below $(1 - \tau)^{1/p}$.

The two panels of Figure 1 show the marginal and average tax rates of the tax schedule (3), for two values of the progressivity parameter: $p = 0.151$, which is calibrated to the rate of progressivity of the US tax code (see Section 3.5) and $p = 0.156$. The second panel shows the

¹In the frictionless model, this can result from assuming that the aggregate production function is linear in the total effective labor input.

tax schedules at the bottom of the income distribution.

Figure 1: Tax schedules



2.2 Individual behavior

I now analyze the effects of taxes on the behavior of individuals in the frictionless model set up in Section 2.1.

Optimal taxable income and value function. Individuals choose their optimal stream of taxable incomes $\{y_t^*\}_{t \geq 0}$ contingent on their stream of productivities $\{\theta_t\}_{t \geq 0}$ in order to maximize the expected present discounted value of their utility. They discount the future at rate ρ and receive the job destruction shock (death) at rate β , after which their new productivity draw θ_0 is independent of their past productivity and decisions. Thus the total discount rate at which future utility is discounted is $(\rho + \beta)$. Thus individuals solve the following problem:

$$\mathcal{V}^*(\theta_0) = \max_{\{y_t^*\}_{t \geq 0}} \mathbb{E}_{\theta_0} \left[\int_0^{\infty} e^{-(\rho+\beta)t} U(\theta_t, y_t^*) dt \right]. \quad (5)$$

The indirect utility, or value function, of an individual with current productivity θ_0 , is denoted $\mathcal{V}^*(\theta_0)$.

The solution to this problem is as follows. At each instant t , the optimal labor income y_t^* is an increasing function of their current productivity θ_t and their net-of-tax rate $(1 - T'(y_t^*))$,

$$\begin{aligned} y_t^* &= (1 - T'(y_t^*))^\varepsilon \theta_t^{1+\varepsilon} = (1 - \tau)^{\frac{\varepsilon}{1+p\varepsilon}} \theta_t^{\frac{1+\varepsilon}{1+p\varepsilon}}, \\ c_t^* &= y_t^* - T(y_t^*) = \frac{1}{1-p} (1 - \tau)^{\frac{1+\varepsilon}{1+p\varepsilon}} \theta_t^{\frac{(1-p)(1+\varepsilon)}{1+p\varepsilon}}. \end{aligned} \quad (6)$$

These income choices show that an individual who becomes more productive finds it optimal to provide more effort, and hence earn a higher income. On the other hand, facing a higher marginal

tax rate (i.e., a lower net-of tax rate) induces the individual to reduce his labor supply. The magnitude of the behavioral response of income to changes in productivity or in the marginal tax rates is measured by the value of the structural parameter ε . In particular, the larger this elasticity, the stronger the decrease in labor income in response to a given tax change, that is, the higher the disincentive effect of taxation. Optimal incomes are decreasing both in the parameter τ and in the rate of progressivity p of the tax schedule. Note that in this model, incomes depend only on the *current* levels of productivity and marginal tax rates; in this sense, individual decisions are static. Note finally that the quasilinearity of the utility function implies that there are no income effects on labor supply, which therefore depends only on the *marginal* tax rates faced by the individual.

Using equation (6), we find that the laws of motion of the taxable and disposable incomes are given by the following geometric Brownian motions:

$$\begin{aligned} d \ln y_t^* &= \mu_y dt + \sigma_y d\mathcal{W}_t, \quad \text{with } \{\mu_y, \sigma_y\} = \frac{1 + \varepsilon}{1 + p\varepsilon} \{\mu_\theta, \sigma_\theta\}, \\ d \ln c_t^* &= \mu_c dt + \sigma_c d\mathcal{W}_t, \quad \text{with } \{\mu_c, \sigma_c\} = (1 - p) \{\mu_y, \sigma_y\}. \end{aligned} \tag{7}$$

Thus the log-taxable income and the log-disposable income processes both have a unit root, which is inherited from the unit root of the log-productivity process. The growth rate of taxable (resp., disposable) income is equal to $g_y = \mu_y + \sigma_y^2/2$ (resp., $g_c = \mu_c + \sigma_c^2/2$). In particular, the taxable income process (7) is a random growth process, so that earning shocks in my model are permanent. A large literature in labor economics estimates income processes of exactly this type, see e.g. Meghir and Pistaferri (2004, 2011).² Note that here, I obtain this income process *endogenously* from the productivity process; it is determined explicitly as a function of the elasticity of labor supply and the rate of progressivity of the tax schedule. A higher elasticity and a lower rate of progressivity lead to a higher volatility of the income process.

I assume that $\rho + \beta - \mu_c - \frac{1}{2}\sigma_c^2 > 0$, which ensures that the individual indirect utility is finite. The present discounted value of the indirect utility $\mathcal{V}^*(y_0^*)$ of an individual with current productivity θ_0 at time zero, and hence taxable income $y_0^* = y_0^*(\theta_0)$ and consumption $c_0^* = \frac{1-\tau}{1-p} (y_0^*)^{1-p}$, is given by:

$$\mathcal{V}^*(y_0^*) = \frac{1}{\rho + \beta - \mu_c - \frac{1}{2}\sigma_c^2} \frac{1 + p\varepsilon}{1 + \varepsilon} c_0^*. \tag{8}$$

Note that the relevant discount rate to compute the individual value function (8) depends on the growth rate of future consumption g_c , and hence is endogenous to taxes. Thus, even though individual behavior is essentially static, i.e., it depends only on the current levels of productivity and taxes that the agent faces, the value of individual welfare takes into account

²Other authors, e.g., Guvenen, Karahan, Ozkan, and Song (2014) argue that different income shock processes (e.g., mixtures of lognormals) better fit the data. From a theoretical viewpoint, the random growth process has the advantage of being extremely tractable, and of generating very naturally Pareto tails for the income distributions.

their forward looking behavior. Moreover, the income distributions will be characterized in the next section as the stationary aggregate outcome of these individual income processes. Therefore, embedding the *purely static* canonical model of taxation (see, e.g., Diamond 1998) into a dynamic environment allows me to treat this framework as a model of *long-run* taxation.

Effects of tax policy on individual income choices. Equations (6) and (7) give closed-form expressions for the individual frictionless choice variables y^*, c^* and their evolution as a function of the parameters (τ, p) of the tax schedule. In response to an infinitesimal perturbation $(d\tau, dp)$ of the initial tax schedule, the first-order change in frictionless taxable income y^* is given by

$$d \ln y^* = \left(\frac{\varepsilon}{1 + p\varepsilon} \right) d \ln (1 - \tau) - \left(\frac{\varepsilon}{1 + p\varepsilon} \ln y^* \right) dp. \quad (9)$$

The interpretation of equation (9) is as follows. The behavioral change in income following a tax increase (both for an increase in τ and in p) is determined by the structural parameter ε . In particular, if the baseline tax system is linear, i.e. $p = 0$, (9) implies immediately that the elasticity of labor income y^* with respect to the net-of tax rate $(1 - \tau)$ is equal to ε . Suppose now that the baseline tax system is progressive or regressive, i.e. $p \neq 0$. Then a change in the marginal tax rate $T'(y^*)$ that an individual faces induces a *direct* reduction of his labor income y^* by the amount ε , by definition of the labor income elasticity. This direct adjustment generates in turn an *indirect* change in the marginal tax rate that the individual faces, due to the non-linearity of the baseline tax schedule. The amount of this change is equal to $d(T'(y)) = T''(y) dy$, and it induces a further labor income adjustment given by the elasticity ε . Thus the total change in income following a perturbation of the net-of-tax rate $(1 - T'(y^*))$ of an individual with income y^* is given by

$$\frac{d \ln y^*}{d \ln (1 - T'(y^*))} = \frac{\varepsilon}{1 + T''(y^*) \frac{y^* \varepsilon}{1 - T'(y^*)}} = \frac{\varepsilon}{1 + p\varepsilon}. \quad (10)$$

Equations (9) and (10) thus show that, from the point of view of individuals, the effect on income of perturbing the parameter $(1 - \tau)$ of the tax schedule by one percent is equivalent to perturbing the net-of tax rate at every income level by one percent. Similarly, the effect of perturbing the parameter p of the tax schedule by dp is equivalent to perturbing the marginal tax rates faced by all individuals by an amount proportional to their log-income, $(\ln y^*) dp$, so that the higher the income, the higher the increase in the marginal tax rate.

Finally, the effects of a perturbation $(d\tau, dp)$ of the initial tax schedule on the drift and volatility of the taxable and disposable income processes are given by:

$$\begin{aligned} \frac{d \ln \{\mu_y, \sigma_y\}}{d \ln (1 - \tau)} &= 0 \quad \text{and} \quad \frac{d \ln \{\mu_y, \sigma_y\}}{dp} = -\frac{\varepsilon}{1 + p\varepsilon} < 0, \\ \frac{d \ln \{\mu_c, \sigma_c\}}{d \ln (1 - \tau)} &= 0 \quad \text{and} \quad \frac{d \ln \{\mu_c, \sigma_c\}}{dp} = -\frac{1}{1 - p} \frac{1 + \varepsilon}{1 + p\varepsilon} < 0. \end{aligned}$$

Therefore, a higher degree of progressivity in the tax schedule leads to a lower drift and volatility of the taxable and disposable income processes. Intuitively, individual income responses following an increase in their productivity are attenuated by the fact that higher incomes pay higher marginal tax rates if the tax schedule is progressive.

2.3 Aggregation

In this section, I derive the income distributions that are generated endogenously by the aggregation of individuals' behavior, characterized in Section 2.2.

Stationary income distributions. Let $x \in \{\theta, y^*, c^*\}$ denote either the productivity θ , the optimal taxable income y^* , or the disposable income c^* . Let also $x_0 \in \{\theta_0, y_0^*, c_0^*\}$ denote the productivity, taxable income, or disposable income at re-entry into the labor force, and denote by $f_{x_0}(\cdot)$ their corresponding (exogenous) distributions. Since θ_0 is log-normally distributed with mean m_θ and variance s_θ^2 , taxable income and consumption at reentry are also log-normally distributed, i.e., $\ln y_0^* \sim \mathcal{N}(m_y, s_y^2)$ and $\ln c_0^* \sim \mathcal{N}(m_c, s_c^2)$. Their mean and variance are given by:

$$\begin{aligned} m_y &= \frac{1+\varepsilon}{1+p\varepsilon}m_\theta + \frac{\varepsilon}{1+p\varepsilon}\ln(1-\tau), \quad \text{and } s_y = \frac{1+\varepsilon}{1+p\varepsilon}s_\theta, \\ m_c &= (1-p)m_y + \ln\left(\frac{1-\tau}{1-p}\right), \quad \text{and } s_c = (1-p)s_y. \end{aligned} \tag{11}$$

The higher the mean and variance (m_θ, s_θ) of the productivity distribution at re-entry, the higher the mean and variance (m_y, s_y) and (m_c, s_c) of the taxable and disposable income distributions at re-entry. As usual, the elasticity of labor income and the parameters of the tax schedule determine the amount by which higher moments of the productivity distribution translate into higher moments of the income distributions. Note that only the mean of the income distribution is affected by a uniform increase in the marginal tax rates (τ) , while the rate of progressivity affects both the mean and the variance. Intuitively, an increase in τ has the same multiplicative effects on all incomes, and hence does not change the variance of the distribution, while an increase in p compresses the distribution by increasing the tax rates by more for higher incomes.

Define the coefficients $r_{1,x} < 0 < r_{2,x}$ for $x \in \{\theta, y^*, c^*\}$ as

$$r_{1,x}, r_{2,x} = \frac{\mu_x}{\sigma_x^2} \pm \sqrt{\left(\frac{\mu_x}{\sigma_x^2}\right)^2 + \frac{2}{\sigma_x^2}\beta}. \tag{12}$$

The following proposition characterizes the productivity and income distributions in the frictionless economy:

Proposition 1. *The distributions of productivity θ , frictionless before-tax income y^* and disposable income c^* converge towards unique stationary distributions. Letting $x \in \{\theta, y^*, c^*\}$, these*

distributions are double-Pareto-lognormal with parameters $(r_{1,x}, r_{2,x}, m_x, s_x^2)$, defined as

$$f_x(x) = \frac{|r_{1,x}|r_{2,x}}{|r_{1,x}| + r_{2,x}} \left\{ e^{\frac{1}{2}r_{1,x}^2 s_x^2 - r_{1,x}m_x} x^{r_{1,x}-1} \Phi\left(\frac{\ln x - m_x}{s_x} + r_{1,x}s_x\right) + e^{\frac{1}{2}r_{2,x}^2 s_x^2 - r_{2,x}m_x} x^{r_{2,x}-1} \Phi\left(\frac{\ln x - m_x}{s_x} + r_{2,x}s_x\right) \right\}. \quad (13)$$

In particular, the densities of productivity θ , desired taxable income y^* , and desired disposable income c^* exhibit power-law behavior in both tails, with Pareto coefficients on the right and left tail respectively given by $(r_{1,\theta}, r_{2,\theta})$, $(r_{1,y}, r_{2,y})$ and $(r_{1,c}, r_{2,c})$, defined in (12). That is,

$$f_x(x) \sim_{x \rightarrow 0} x^{r_{2,x}-1} \quad \text{and} \quad f_x(x) \sim_{x \rightarrow \infty} x^{r_{1,x}-1}. \quad (14)$$

Proof. See Appendix. □

Proposition 1 characterizes the productivity and income distributions in closed form. Several authors (e.g., Reed 2003, Reed and Jorgensen 2004, Toda 2012) argue that the double-Pareto-lognormal distribution fit very well the actual income distributions observed empirically. In my model, these distributions are generated *endogenously* as a function of the parameters of the tax schedule. Importantly, these distributions have Pareto tails, which is one of the most established stylized fact about the actual income distributions (e.g., Nirei and Souma, 2007) and plays an important role in the theory of taxation (Saez 2001). The corresponding *Pareto coefficients* $r_{1,y}, r_{2,y}$ are given by (12) as a function of taxes. Note that these characteristics of the *aggregate* distributions are determined by the *individual* idiosyncratic income processes, i.e., the drift and volatility (μ_y, σ_y) . It is well known that the aggregation of random growth process of the form (7) generates very naturally stationary distributions that have Pareto tails (Gabaix 2009). In the bulk of the distribution, i.e. away from the tails, the distribution is approximately lognormal, a feature inherited from the lognormal density of productivities at re-entry into the labor force.

In particular, the parameter $|r_{1,y}|$ is the Pareto coefficient of the right tail of the income distribution and determines its thinness. The higher this coefficient (in absolute value), the thinner the tail, the more equal the income distribution. A higher drift μ_y of individual income leads to a more unequal the distribution, i.e., a smaller value of $|r_{1,y}|$. A higher volatility σ_y leads to a more unequal distribution. Finally, a higher rate of job destruction or death β induces a less unequal income distribution. The following relationships relate the Pareto coefficients of the taxable and disposable income distributions to those of the productivity distribution:

$$\{r_{1,y}, r_{2,y}\} = \frac{1+p\varepsilon}{1+\varepsilon} \{r_{1,\theta}, r_{2,\theta}\} \quad \text{and} \quad \{r_{1,c}, r_{2,c}\} = \frac{1}{1-p} \frac{1+p\varepsilon}{1+\varepsilon} \{r_{1,\theta}, r_{2,\theta}\}.$$

That is, the Pareto coefficients of the taxable and disposable income distributions are inherited from those of the productivity distribution, and are a function of the labor income elasticity and

the rate of progressivity of the tax schedule. The distribution of desired taxable income is more unequal (thicker tail) than the productivity distribution, because productivity differences are amplified by the positive labor supply elasticity ε . Moreover, a higher degree of progressivity increases the values of these parameters, i.e., makes the distributions less unequal. The Pareto parameters satisfy $r_{1,y} \leq r_{1,c} < r_{1,\theta}$ if $p \geq 0$ and $r_{1,c} \leq r_{1,y} < r_{1,\theta}$ if $p \leq 0$, where the inequalities are strict if $p \neq 0$. The distribution of frictionless disposable income is less unequal (thinner tail) than the distribution of desired taxable income if and only if the tax schedule is progressive, $p > 0$. Although both distributions are affected, the effect of a higher rate of progressivity on the after-tax income distribution is stronger than on the pre-tax distribution. Note that the parameter τ of the tax schedule does not affect the Pareto tails. This is because the Pareto coefficients of the distribution of income $x \in \{y, c\}$ depend on the ratio between the average income above a threshold \bar{x} , $\mathbb{E}_{x \geq \bar{x}}[x]$, and the threshold \bar{x} (as $\bar{x} \rightarrow \infty$). Changing the scaling parameter τ has the same multiplicative effect on the incomes of each individual, so that the ratio $\mathbb{E}_{x \geq \bar{x}}[x] / \bar{x}$ is unaffected.

Effects of tax policy on the aggregate income distributions. An increase in the parameter τ , which corresponds to a uniform increase in the marginal tax rates by the same amount at every income level, only reduces the means m_y, m_c of the lognormal bulks of the taxable and disposable income distributions. An increase in the rate of progressivity p reduces both the mean and the variance of the lognormal bulk, with

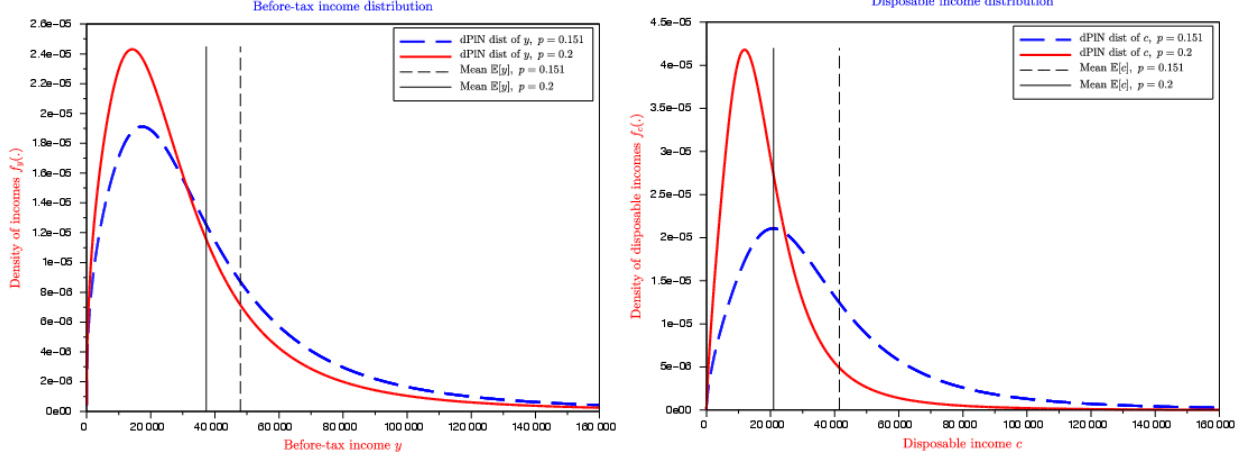
$$\frac{dm_y}{d \ln(1 - \tau)} = \frac{\varepsilon}{1 + p\varepsilon}, \quad \text{and} \quad \frac{d \ln \{m_y, s_y\}}{dp} = -\frac{\varepsilon}{1 + p\varepsilon}.$$

Moreover, an increase in progressivity p increases the Pareto coefficients of the income distributions, i.e., both the before-tax and the after-tax income distributions have a thinner tail if the degree of progressivity of the tax schedule is higher. We find:

$$\frac{d \ln \{r_{1,y}, r_{2,y}\}}{dp} = \frac{\varepsilon}{1 + p\varepsilon}, \quad \text{and} \quad \frac{d \ln \{r_{1,c}, r_{2,c}\}}{dp} = \frac{1 + \varepsilon}{1 + p\varepsilon}.$$

The two panels of Figure 2 show the aggregate taxable and disposable income distributions calibrated to the U.S. economy ($p = 0.151$, see Section 3.5), and the effect of an increase in the progressivity parameter p (from 0.151 to 0.2) on these distributions. That is, the net-of tax rate decreases by 0.20 percent rather than 0.15 percent when income increases by 1 percent. The figure shows that effect of higher progressivity on the thinness of the tail of the after-tax income distribution is much stronger than on the pre-tax income distribution. That is, inequality in after-tax incomes reduces by more than inequality in pre-tax incomes.

Figure 2: Aggregate taxable income and disposable income distributions



3 Long-run effects of taxation on social welfare

In this section, I analyze the effects of taxes on long-run social welfare. I suppose that the distributions of productivities and incomes have converged to their stationary distributions characterized in Section 2.3, and study the effects of changes in taxes. I first define the objective of the government and the marginal social welfare weights. Those are endogenous variables, measured in monetary units, which characterize the redistributive tastes of the government and are used as sufficient statistics to characterize the effects of tax policy on social welfare.

3.1 Government and social welfare

The government chooses the tax schedule $T = \{\tau, p\}$. I consider two distinct long-run social welfare criteria, \mathcal{W}_{pg} and \mathcal{W}_r , corresponding respectively to (i) a public good provision problem, and (ii) a redistribution problem. Let $\mathcal{R}(T)$ denote the long-run present discounted value, at rate ρ , of tax revenue given the tax schedule T , defined as

$$\mathcal{R}(T) = \rho^{-1} \int_0^{\infty} T(y) dF_y(y), \quad (15)$$

where $f_y(\cdot)$ (resp., F_y) is the stationary density (resp., c.d.f.) of taxable incomes y in the economy given the tax schedule T .

First, I define social welfare in the *public good provision problem* as

$$\mathcal{W}_{\text{pg}}(T) = \frac{1}{\lambda^*} \left[\int_0^{\infty} \mathcal{G}(v^*(y)) dF_y(y) + \mathcal{R}(\mathcal{R}(T)) \right]. \quad (16)$$

The first term in the brackets of (16) is the social objective, equal to the sum of individual indirect utilities. The function $\mathcal{G} : \mathbb{R} \rightarrow \mathbb{R}$ denotes the social welfare function, defined over lifetime

indirect utilities $\mathcal{V}^*(y)$ of individuals and assumed continuously differentiable, increasing, and concave. The more concave the social welfare function \mathcal{G} , the more the planner wants to redistribute from individuals with high income to those with low income. The function $\mathcal{R} : \mathbb{R} \rightarrow \mathbb{R}$ denotes the value of public goods that the government can provide with the tax revenue $\mathcal{R}(T)$ collected, and is assumed continuously differentiable and increasing. Finally, the real number $\lambda^* \equiv \mathcal{R}'(\mathcal{R})$ denotes the (exogenous) marginal value of public funds. Normalizing social welfare by λ^* gives us a money metric to evaluate the welfare effects of tax changes. I assume that the rate of progressivity p in this problem is fixed, and the government can only change the marginal tax rates uniformly to provide more or less of the public good; that is, the only available (distortionary) tax instrument is the parameter τ . In the sequel, I derive the effects on government tax revenue \mathcal{R} , and hence on social welfare \mathcal{W}_{pg} , of perturbing the parameter τ .

Second, I define social welfare in the *redistribution problem* as

$$\mathcal{W}_r(T) = \frac{1}{\lambda^*} \int_0^\infty \mathcal{G}(\mathcal{V}^*(y)) dF_y(y), \text{ with } \mathcal{R}(T) \geq \bar{R}. \quad (17)$$

In this problem, the government can change the rate of progressivity p to increase the social objective, subject to the constraint that the total tax revenue levied must be larger than or equal to an exogenous revenue requirement \bar{R} . Imposing $\mathcal{R}(T) = \bar{R}$ pins down the parameter τ for a given progressivity p . Importantly, λ^* denotes the (endogenous) marginal value of public funds in this problem, which I define formally in the next section; this implies that the social welfare criterion \mathcal{W}_r is expressed in monetary units (dollars).

In the redistribution problem, the government can use both tax instruments, τ and p , subject to satisfying its revenue requirement. In the sequel, I consider two possible tax reforms of a given baseline tax system. First, I derive the effects on government tax revenue \mathcal{R} of changing the marginal tax rates (τ) or the rate of progressivity (p). Second, I compute the effects on social welfare \mathcal{W}_r of *budget-neutral* changes in the degree of progressivity of the tax schedule, i.e., where the parameter τ adjusts so that total tax revenue is unchanged and equal to \bar{R} . In this case, the marginal value of public funds λ is obtained by imposing that a small increase in τ (by $d\tau \rightarrow 0$) has no effect on social welfare.

Note finally that in my model individuals are risk-neutral and the government weighs their utilities using a concave social welfare function. The results of this paper would be similar if instead individuals were risk-averse with Greenwood, Hercowitz, and Human (1988) preferences, savings and borrowings were banned, and the government was utilitarian. In my formulation, all the concavity (that of the individual utility and that of the planner's redistributive objective) is summarized in the social welfare function.

3.2 Marginal social welfare weights

The effects of tax reforms on the social welfare criteria \mathcal{W}_{pg} and \mathcal{W}_r can be characterized using the notion of marginal social welfare weights, introduced by Mirrlees (1971) (see also Saez

and Stantcheva 2014 for a generalization). Intuitively, the weight at income y represents the increase in social welfare, expressed in terms of public revenue, of distributing uniformly among individuals who earn income y an additional dollar of consumption. I now define formally these welfare weights in the context of my model, for a given social welfare function \mathcal{G} . In order to make the welfare comparisons meaningful and obtain a monetary measure of welfare gains, these weights must be weighted by the marginal value of public funds λ^* .

Equation (8) shows that in the frictionless model of Section 2, the value function of an individual who earns income y_0 , $\mathcal{V}^*(y_0)$, is proportional to the current disposable income $c_0 = \frac{1-\tau}{1-p}y_0^{1-p}$. First, consider the effect on this individual's welfare of giving him an additional consumption stream $\{\hat{c}_t\}_{t \geq 0}$ which evolves stochastically over time according to the same process (μ_c, σ_c) as his frictionless disposable income c_t^* . The present discounted value of his utility given this additional consumption stream is given by:

$$\mathcal{V}^*(y_0; \hat{c}) = \mathcal{V}^*(y_0) + \mathbb{E}_0 \left[\int_0^\infty e^{-(\rho+\beta)t} \hat{c}_t dt \right] = \mathcal{V}^*(y_0) + \frac{\hat{c}_0}{\rho + \beta - \mu_c - \frac{1}{2}\sigma_c^2}. \quad (18)$$

That is, the change in the agent's value function is equal to the present value of the consumption stream, which in turn is equal to the additional consumption at time 0, \hat{c}_0 , discounted at a rate which is increasing in the rate of growth of the consumption stream.

The marginal social welfare weight $\gamma_{\hat{c}}^*(y_0)$ at income level y_0 is then defined as the social value of giving this marginal consumption stream $\{\hat{c}_t\}_{t \geq 0}$ uniformly to all the individuals with income y_0 at time 0. We express this social value in terms of public revenue by normalizing it by the marginal value of public funds λ^* . That is,

$$\frac{\gamma_{\hat{c}}^*(y_0)}{\lambda^*} \equiv \frac{1}{\lambda^*} \left. \frac{d\mathcal{V}^*(y_0; \hat{c})}{d\hat{c}_0} \right|_{\hat{c}_0=0} \mathcal{G}'(\mathcal{V}^*(y_0)) = \frac{1}{\lambda^*} \frac{\mathcal{G}'(\mathcal{V}^*(y_0))}{\rho + \beta - \mu_c - \frac{1}{2}\sigma_c^2}. \quad (19)$$

It is proportional to the marginal increase in social welfare induced by the higher individual utility.

Second, consider the effect on an individual's welfare of increasing the growth rate $g_c = \mu_c + \frac{1}{2}\sigma_c^2$ of his future consumption process by an amount dg_c . As (8) shows, this is equivalent to a decrease in his discount rate ρ . This induces a change in the present discounted value of his utility given, to a first order in dg_c as $dg_c \rightarrow 0$, by:

$$\frac{d\mathcal{V}^*(y_0)}{dg_c} = \frac{1+p\varepsilon}{1+\varepsilon} \frac{c_0}{\left(\rho + \beta - \mu_c - \frac{1}{2}\sigma_c^2\right)^2}. \quad (20)$$

The marginal social welfare weight $\gamma_\rho^*(y_0)$ at income level y_0 is then defined as the social value of this marginal increase in g_c for individuals with income y_0 at time 0. That is,

$$\frac{\gamma_\rho^*(y_0)}{\lambda^*} \equiv \frac{1}{\lambda^*} \frac{d\mathcal{V}^*(y_0)}{dg_c} \mathcal{G}'(\mathcal{V}^*(y_0)) = \frac{1}{\lambda^*} \frac{1+p\varepsilon}{1+\varepsilon} \frac{\mathcal{G}'(\mathcal{V}^*(y_0))}{\left(\rho + \beta - \mu_c - \frac{1}{2}\sigma_c^2\right)^2} c_0. \quad (21)$$

These welfare weights are new to the taxation literature, and come from the embedding of the standard static model into a dynamic environment, which implies that tax policy affects not only the current consumption of individuals, but also its growth rate and hence their future prospects.

These welfare weights summarize the government redistributive tastes, and are the key to measure the changes in social welfare brought about by tax changes. It is important to note that, keeping with the Mirrleesian tradition in public finance, these weights are *not* exogenous, as they depend on the current income of the individual.

3.3 Marginal tax rates and provision of public goods

In this section, I analyze the effects of tax policy in the public good provision problem (16), in which the government can change the parameter τ of the baseline tax schedule $T = \{\tau, p\}$ in order to raise or lower the amount of public good that is provided through tax revenue. Specifically, I consider a local perturbation of τ of the baseline tax system by the small amount $d\tau$, and characterize the first-order effects of this tax reform on tax revenue \mathcal{R} or on the social welfare criterion \mathcal{W}_{pg} as $d\tau \rightarrow 0$.

The tax liability levied at the income level y after the tax reform is implemented is given, to a first order in $d\tau \rightarrow 0$, by

$$\tilde{T}(y) = y - \frac{1 - \tau - d\tau}{1 - p} y^{1-p} = T(y) + \mathbf{t}_\tau(y) d\tau,$$

where $\mathbf{t}_\tau(y) = \frac{y^{1-p}}{1-p}$. Thus, the tax liability at income y changes by $\mathbf{t}_\tau(y) d\tau$, and the marginal tax rate changes by $\mathbf{t}'_\tau(y) d\tau$. I first analyze the effects of the tax reform on the government's tax revenue \mathcal{R} , which I denote by P_τ^* . For ease of interpretation I normalize this effect by the statutory increase in tax revenue, namely,

$$P_\tau^* \equiv \frac{1}{\mathbb{E}[\mathbf{t}_\tau(y)]} \left. \frac{d\mathcal{R}}{d\tau} \right|_{(\tau,p)}.$$

Thus P_τ^* measures the *actual* change in government tax revenue of a one-dollar *statutory* increase in tax revenue due to a uniform increase in the marginal tax rates. I then analyze the effects of the tax reform on social welfare \mathcal{W}_{pg} , which are given by:

$$\Gamma_{\text{pg}}^* \equiv \frac{1}{\mathbb{E}[\mathbf{t}_\tau(y)]} \left. \frac{d\mathcal{W}_{\text{pg}}}{d\tau} \right|_{(\tau,p)}.$$

Thus Γ_{pg}^* measures the change in social welfare, expressed in monetary units, of a one dollar statutory increase in tax revenue (and hence in the amount of public good provided) through a uniform increase in the marginal tax rates. I show:

Proposition 2. *The long-run effect of a uniform increase in marginal tax rates on the government's tax revenue is given by*

$$P_{\tau}^* = 1 - \int_0^{\infty} \frac{T'(y)}{1 - T'(y)} e_{y,1-T'}^* \frac{y \mathbf{t}'_{\tau}(y)}{\mathbb{E}[\mathbf{t}_{\tau}(y)]} f_y(y) dy, \quad (22)$$

where $e_{y,1-T'}^* = \frac{\varepsilon}{1+p\varepsilon}$ is the elasticity of frictionless income with respect to the net-of-tax rate, and $f_y(y)$ is the stationary density of incomes given the baseline tax system $T = \{\tau, p\}$. The long-run effect of a uniform increase in marginal tax rates on social welfare is given by

$$\Gamma_{\text{pg}}^* = P_{\tau}^* - \mathbb{E} \left[\frac{\gamma_{\hat{c}}^*(y)}{\lambda^*} \frac{\mathbf{t}_{\tau}(y)}{\mathbb{E}[\mathbf{t}_{\tau}(y)]} \right], \quad (23)$$

where P_{pg}^* is the revenue effect characterized in (22), $\gamma_{\hat{c}}^*(y)$ is the marginal social welfare weight at income level y , and λ^* is the marginal value of public funds, defined by $\lambda^* = \mathcal{R}'(\mathcal{R})$.

Proof. See Appendix. □

The interpretation of the formulas of Proposition 2 is as follows. The first term in the right hand side of (23) is the *mechanical effect* of the perturbation, i.e., the statutory increase in government revenue absent behavioral responses. This is equal to one (dollar) by construction of the perturbation. The second term in the right hand side of (23) is the *behavioral effect* of the perturbation. The increase $\mathbf{t}'_{\tau}(y) d\tau$ in the marginal tax rate of an individual with income y generated by the perturbation induces him to decrease his taxable income by $\frac{y}{1-T'(y)} \frac{\varepsilon}{1+p\varepsilon} \mathbf{t}'_{\tau}(y) d\tau$, since the labor income elasticity (including the feedback effect on the marginal tax rates due to the non-linearity of the tax schedule) is given by (10) and equal to $\frac{\varepsilon}{1+p\varepsilon}$. This behavioral income response generates a loss in government revenue proportional to the marginal tax rate $T'(y)$. Finally, summing over individuals using the density of incomes $f_y(\cdot)$ in the baseline tax system $\{\tau, p\}$ yields equation (23).

The effects of the perturbation on social welfare, characterized by formula (23), are equal to these revenue effects P_{τ}^* , plus an additional term, which captures the *welfare effect* of the perturbation. An increase in the tax liability of individual y by $\mathbf{t}_{\tau}(y) d\tau$ reduces his utility and hence social welfare by $(\lambda^*)^{-1} \gamma_{\hat{c}}^*(y) \times \mathbf{t}_{\tau}(y) d\tau$, by construction of the marginal social welfare weight $\gamma_{\hat{c}}^*$ at income level y . These two equations have formally the same structure as the “sufficient statistic” formulas derived by Saez (2001) and Golosov, Tsyvinski, and Werquin (2014).

Formula (23) can be used to compute the effects on the social welfare criterion \mathcal{W}_{pg} of *local tax reforms* (specifically, small uniform changes in the marginal tax rates) of any given baseline tax schedule, e.g., the U.S. tax code. For this purpose, the density of incomes $f_y(\cdot)$ that is used to compute the welfare and behavioral effects of the perturbation is the U.S. income distribution. Formula (23) can also be used to characterize the *optimal tax schedule*. Specifically, consider the optimal choice of constant marginal tax rate τ^* in the public good provision problem (16), with

$p = 0$. The optimum is such that no tax reform yields a strictly positive welfare gain. Imposing $\Gamma_{\text{pg}} = 0$ in (23) yields:

Corollary 1. *The optimal tax rate τ^* in the public good provision problem is given by:*

$$\frac{\tau^*}{1 - \tau^*} = \frac{1}{\varepsilon} \left(1 - \mathbb{E} \left[\frac{\gamma_{\hat{c}}^*(y)}{\lambda^*} \frac{y}{\mathbb{E}[y]} \right] \right), \quad (24)$$

where $\mathbb{E}[y] = \int_0^\infty y f_y(y) dy$ is the average labor income given the tax rate τ^* .

Proof. See Appendix. □

Corollary 1 shows that the optimal tax rate τ^* is decreasing in the labor income elasticity ε , and is specifically set according to the *inverse elasticity rule* standard in the public finance literature. Moreover, τ^* is decreasing in the income-weighted average of the marginal social welfare weights. This term balances the social welfare cost of increasing the tax rate τ (by $d\tau$, say), making individual y lose income $y d\tau$, with the benefit of increasing the provision of public goods through higher tax revenue, measured by the marginal value of public funds λ^* . Importantly, in the optimal tax formula (24) the marginal social welfare weights $\gamma_{\hat{c}}^*(y)$ and the income distribution $f_y(\cdot)$ are endogenous to the tax rate τ^* , and cannot be computed using, e.g., the current U.S. income distribution. Note however that the closed-form expressions I obtained in (13) for the before- and after-tax income distributions allow me to compute explicitly them at the optimal schedule.

3.4 Progressivity and redistribution

In this section, I analyze the redistribution problem, in which the government can choose the parameter τ and the rate of progressivity p of the baseline tax schedule subject to satisfying a budget constraint, and where the social welfare criterion is given by (16).

The first tax reform I consider is the following. I derive the effects on tax revenue \mathcal{R} of perturbing p by the small amount dp . The post-perturbation tax liability paid at the income level y is given by, to a first order in $dp \rightarrow 0$,

$$\tilde{T}(y) = y - \frac{1 - \tau}{1 - p - dp} y^{1-p-dp} = T(y) + \mathbf{t}_p(y) dp,$$

where $\mathbf{t}_p(y) = \frac{1-\tau}{1-p} y^{1-p} \left(\ln y - \frac{1}{1-p} \right)$. Thus, the tax liability at income y changes by $\mathbf{t}_p(y) dp$, and the marginal tax rate changes by $\mathbf{t}'_p(y) dp$. I compute the first order effect in dp of this perturbation on tax revenue, normalized by the statutory increase in tax revenue. That is,

$$P_p^* \equiv \frac{1}{\mathbb{E}[\mathbf{t}_p(y)]} \frac{d\mathcal{R}}{dp} \Big|_{(\tau,p)}.$$

Thus P_p^* measures the *actual* increase in tax revenue of a one dollar *statutory* increase in revenue through an increase in the rate of progressivity p . The effects on tax revenue of a change in the parameter τ are identical to those derived in Proposition 2, that is, equal to P_τ^* .

The second tax reform I consider is a *budget-neutral* change in p by dp , that is, combined with a change in τ by $d\tau$ so that total tax revenue \mathcal{R} remains unchanged (at \bar{R} , say). Thus, the size of the perturbation $d\tau$ is pinned down by

$$\left. \frac{d\mathcal{R}}{dp} \right|_{(\tau,p)} dp + \left. \frac{d\mathcal{R}}{d\tau} \right|_{(\tau,p)} d\tau = 0.$$

I then compute the first order effect in dp of this perturbation on social welfare, that is,

$$\Gamma_p^* \equiv \left. \frac{d\mathcal{W}_r}{dp} \right|_{(\tau,p)} - \frac{d\mathcal{R}/dp}{d\mathcal{R}/d\tau} \left. \frac{d\mathcal{W}_r}{d\tau} \right|_{(\tau,p)}.$$

Thus Γ_p^* measures the change in social welfare, *expressed in monetary units* (as \mathcal{W}_r is scaled by the marginal value of public funds λ^*), of a revenue-neutral increase in the rate of progressivity p . I show:

Proposition 3. *In the frictionless model, the effect on tax revenue of a change in the parameter τ is given by P_τ^* , and the effect on tax revenue of a change in the rate of progressivity p is given by*

$$P_p^* = 1 - \int_0^\infty \frac{T'(y)}{1 - T'(y)} e_{y,1-T'}^* \frac{y \mathbf{t}'_p(y)}{\mathbb{E}[\mathbf{t}_p(y)]} f_y(y) dy. \quad (25)$$

The effect on social welfare of a budget-neutral change in the rate of progressivity p is given by

$$\Gamma_r^* = P_p^* - \mathbb{E} \left[\frac{\gamma_{\hat{c}}^*(y)}{\lambda^*} \frac{\mathbf{t}_p(y)}{\mathbb{E}[\mathbf{t}_p(y)]} - \frac{\gamma_\rho^*(y)}{\lambda^*} \frac{dg_c/dp}{\mathbb{E}[\mathbf{t}_p(y)]} \right], \quad (26)$$

where the marginal value of public funds λ^* is equal to:

$$\lambda^* = \frac{1}{P_\tau^*} \mathbb{E} \left[\gamma_{\hat{c}}^*(y) \frac{\mathbf{t}_\tau(y)}{\mathbb{E}[\mathbf{t}_\tau(y)]} \right]. \quad (27)$$

Proof. See Appendix. □

Proposition 3 shows that the effects on tax revenue and social welfare of an increase in the rate of progressivity are determined by sufficient statistic formulas similar to (22) and (23), except that the change in the tax liability and the marginal tax rate corresponding to this perturbation are determined by the functions $\mathbf{t}_p(y)$ and $\mathbf{t}'_p(y)$ rather than $\mathbf{t}_\tau(y)$ and $\mathbf{t}'_\tau(y)$. Otherwise, the interpretation of the formulas is similar to that discussed in the context of Proposition 2.

There are two important differences, however. First, the welfare loss associated with an increase in p (the second term in the right hand side of (26)) includes an additional term, which is determined by the marginal social welfare weights $\gamma_\rho^*(y)$. This novel term in the literature

captures the effect of progressivity on the present discounted value of individual utility through its effect on the growth rate of future consumption g_c . As discussed in Section 2.2, in this model tax policy affects not only the current level consumption, but also curbs the drift and volatility of the consumption process, which the individual takes into account when calculating the present value of his utility. The second important difference between the redistribution problem of Proposition 3 and the public good provision problem of Proposition 2 is that in the former case, the marginal value of public funds λ^* is not exogenous as in (23), but is determined as a function of the marginal social welfare weights $\gamma_c^*(y)$. Specifically, λ^* is pinned down by imposing that perturbing τ has no effect on social welfare, so that the government budget constraint in the redistribution problem is not slack and the parameter τ is set optimally given the rate of progressivity p . Thus, as far as the welfare criterion \mathscr{W}_r is concerned, the marginal value of public funds λ^* is equal to the social value of redistributing a dollar of tax revenue through a decrease in τ , i.e., through a uniform decrease in the marginal tax rates.

3.5 Calibration of the model and quantitative analysis

Calibration of the model’s parameters. I now turn to the calibration of the frictionless model of Sections 2 and 3. First, I calibrate the rate of progressivity p of the tax schedule in the US using the empirical estimate from PSID data of Heathcote, Storesletten, and Violante (2014): $p = 0.151$. This means that a 1 percent increase in income leads to a 0.15 percent decrease in the net-of-tax rate (see equation (4)).

There is substantial controversy in the literature about the value of the taxable income elasticity ε ; the micro literature typically finds values lower than 0.3, while the macro literature and some structural estimates find it to be closer to 1 (see Saez, Slemrod, and Giertz 2012, and Keane and Rogerson 2012, for an overview of the two strands). Chetty (2012) estimates the structural parameter (Hicksian intensive margin elasticity) $\varepsilon = 0.33$ using a meta analysis of micro and macro studies and allowing for adjustment frictions. In my baseline calibration I take $\varepsilon = 0.5$, which is a mid-range estimate of the empirical estimates and close to the estimate of Gruber and Saez (2002). I discuss below how my main results are affected by the value of the elasticity ε .

I calibrate the Pareto coefficients of the *observed* distribution of incomes, $r_{y,1}$ and $r_{y,2}$. The coefficient of the right tail, $|r_{y,1}|$, is well known: it varies around 2 and has been decreasing (the tail of the distribution has become thicker, i.e. more unequal) in the past few decades. I take $(r_{y,1}, r_{y,2}) = (-1.8, 1.4)$, see e.g. Reed (2003), Reed and Jorgensen (2004).

The mean m_y and variance s_y of the lognormal “bulk” of the income distribution are calibrated using the mean and variance of the actual distribution of log-incomes,

$$\mathbb{E}[\ln y] = m_y - \frac{1}{r_{y,1}} - \frac{1}{r_{y,2}} = 10.3, \quad \text{and} \quad \mathbb{V}[\ln y] = s_y^2 + \frac{1}{r_{y,1}^2} + \frac{1}{r_{y,2}^2} = 1.$$

Using the calibrated values of the Pareto coefficients, I obtain $(m_y, s_y) = (10.46, 0.43)$.

There is a large literature estimating log-income dynamics that follow a geometric random walk, that is equation (7), see e.g. Meghir and Pistaferri (2004, 2011). The volatility of idiosyncratic income risk σ_y^2 in my model corresponds to the variance of the permanent component of the individual log-income process in this literature. I calibrate $\sigma_y^2 = 0.01$ (see also Jones and Kim 2014, for an estimate in a model similar to mine and further references to the empirical literature).

All the other parameters of my model are then pinned down. In particular, note that the *cross-sectional* income distribution in the economy, specifically the values of the two Pareto coefficients at the tails, allows us to infer information about the *time series* of individual income, since these Pareto tails are generated by the underlying random growth process for income. Specifically, we have

$$r_{y,1} + r_{y,2} = \frac{2\mu_y}{\sigma_y^2}, \quad \text{and} \quad r_{y,1}r_{y,2} = -\frac{2\beta}{\sigma_y^2},$$

which pin down the drift μ_y and the destruction rate β . I take a discount rate ρ so that $(1 + \rho + \beta)^{-1} = 0.95$. Note that this implies a (slightly) negative drift μ_y of log-incomes; however the growth rate of income, $g_y = \mu_y + \sigma_y^2/2$, is positive.

Finally, the parameters of the individual productivity and consumption processes, $(\mu_\theta, \sigma_\theta)$ and (μ_c, σ_c) , and those of the productivity and consumption distributions, $(m_\theta, s_\theta, r_{1,\theta}, r_{2,\theta})$ and $(m_c, s_c, r_{1,c}, r_{2,c})$ are then obtained by equations (7), (11), and (12).

In the sequel, I compute the effects of a tax change by keeping the parameters of the exogenous *productivity* process $(\mu_\theta, \sigma_\theta, m_\theta, s_\theta)$ *constant*, and update the parameters of the endogenous income distribution.

Effects of tax reforms. I now illustrate the usefulness of formulas (22,23) and (25,26) to measure the effects of tax policy on the economy. These formulas can be used to compute the effects on tax revenue and social welfare of *local tax reforms* (specifically, small uniform changes in the marginal tax rates, or budget-neutral increases in the rate of progressivity) of any given baseline tax schedule, e.g., the US tax code. Importantly, besides the marginal social welfare weights which summarize the redistributive tastes of the government, they are expressed in terms of parameters that are potentially observable empirically: the labor income elasticity, the distribution of incomes in the economy, and the parameters of the baseline tax schedule. The advantage of using (23) to compute the effects of *locally reforming* the US tax code is that these “sufficient statistics” are estimated given the *baseline* tax system (and are not, for example, those that would be relevant for the characterization of the optimal tax system). I thus use the values calibrated in the previous paragraph to compute the effects of reforming the current U.S. tax code.

Rather than taking a stand on the redistributive tastes of the government, i.e., the social welfare function \mathcal{G} and the value of public goods \mathcal{R} , I show how applying (22) and (25) leads

to straightforward back-of-the-envelope calculations of the revenue effects of the corresponding tax reforms. We find easily

$$P_{\tau}^* = \frac{1 + \varepsilon}{1 + p\varepsilon} - \frac{\varepsilon}{1 + p\varepsilon} \frac{\mathbb{E}[y]}{\mathbb{E}[c]} = 0.84.$$

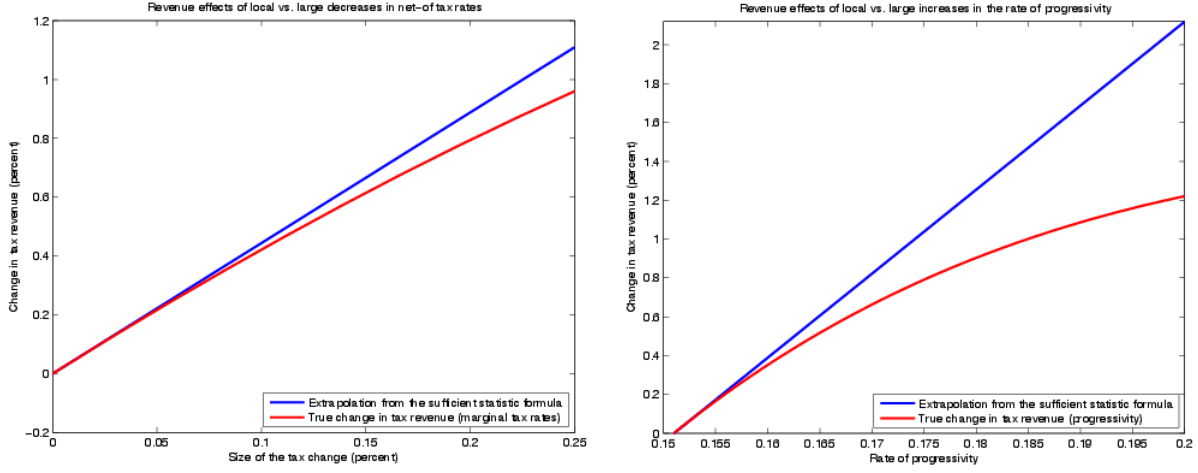
This calculation shows that increasing the *statutory* tax revenue by \$1 through a uniform increase in the marginal tax rates, would yield an *actual* increase in government revenue by ¢84, once the individual behavioral responses have been taken into account. If the elasticity of labor income is larger, $\varepsilon = 1$, then this revenue gain falls to ¢70.6, as the individual behavioral response to higher taxes is stronger. Similarly, raising the progressivity of the US tax code in order to raise a \$1 statutory increase in tax revenue yields

$$P_p^* = 1 - \frac{\varepsilon(1-p)\mathbb{E}[y \ln y] - (1-p)\mathbb{E}[c \ln y]}{1 + p\varepsilon((1-p)\mathbb{E}[c \ln y] - \mathbb{E}[c])} = 0.84,$$

that is, a ¢84 increase in actual tax revenue. If the elasticity of labor income is $\varepsilon = 1$, this revenue gain falls to ¢70.2.

Equations (22) to (26) are thus useful to compute the (revenue or welfare) effects of “small” tax reforms of the current tax schedule. One question is whether this formula provides a good approximation of the effects of *large* tax reforms. Figure 3 shows the revenue gains of tax reforms given by the sufficient statistic formulas, extrapolated to large increases in the tax rates, along with the true changes, in which the deformation of the income distribution is taken into account. The first panel shows the effect of raising the marginal tax rates uniformly (increasing τ), and the second panel shows the effect of raising the rate of progressivity p . An increase in τ by $\frac{\Delta\tau}{1-\tau} = 20\%$ (on the horizontal axis of the first panel) induces an increase in the marginal tax rate at income $\$50.10^3$ (resp., $\$100.10^3$, $\$500.10^3$) from 21.9% (resp., 29.7%, 44.9%) to 37.5% (resp., 43.8%, 55.9%). An increase in p from 0.151 to 0.165 (on the horizontal axis of the second panel) induces an increase in the marginal tax rate at income $\$50.10^3$ (resp., $\$100.10^3$, $\$500.10^3$) from 21.9% (resp., 29.7%, 44.9%) to 46.3% (resp., 52.1%, 63.3%).

Figure 3: Evaluation of the sufficient statistics approach: revenue effects of large tax reforms



4 Labor income adjustment frictions

I now introduce frictions to the adjustment of labor income into the model set up in Section 2. I first set up the environment and characterize individual behavior in this model (Section 4.1), then describe the aggregation (Section 4.2).

4.1 Individual behavior

Adjustment costs. Consider an individual with current productivity θ , and hence optimal frictionless taxable and disposable labor incomes y^* and c^* , respectively. Due to the presence of adjustment frictions, his *actual* current taxable income y may differ from his current *desired* labor income y^* , that is the income he would choose if he could adjust at no cost, which I characterized in Section 6. I assume that in order to adjust his taxable income from y to $y' \neq y$, he incurs a fixed cost κ proportional to his current (frictionless) disposable income $c^*(\theta)$, i.e.,

$$\kappa(\theta) = \kappa \times c^*(\theta),$$

where $\kappa > 0$ denotes the exogenous fixed cost per unit of (frictionless) disposable income. κ can be interpreted as the opportunity cost of time spent searching for a new job, equal to his after-tax income.³ There is moreover an exogenous Poisson arrival of costless adjustment opportunities at rate q .

Intuitively, consider an individual who works in a job with income or hours of work y . If this individual becomes more productive (higher θ), he would like to work more and earn a higher income, which he would do in the frictionless model analyzed in Section 2. In the frictional world, however, to do so he has to find a new job that allows him to work and earn more, and

³I assume for simplicity that the fixed cost is proportional to the frictionless, rather than the actual, disposable income. This is only for tractability.

search is costly. He thus remains in his current occupation, but because he has become more productive, he simply provides less effort to produce the same amount and earn the same income. If y is interpreted as hours of work, he stays the same number of hours at work, but spends some of these hours idle, i.e., not providing (costly) effort e .⁴ He decides to pay the search cost and adjust his hours and income y upward only when his productivity becomes much higher, so that he spends most of his time resting in his current job. He then finds a new and more challenging occupation with income $y' > y$, and so on. Conversely, if the individual becomes less productive (e.g., if his health is declining, or he has a child), he will need to provide more effort during his hours of work to earn the same income, until he becomes so unproductive that he decides to switch jobs and move to a less challenging occupation with $y' < y$. Empirically, the variance of incomes for individuals who change jobs is typically three times or more larger than the variance of incomes for those who stay in their job; in my model, this variance is equal to zero for the stayers, and is positive for the movers. Finally, the individual may randomly and exogenously receive adjustment opportunities (i.e., job offers) at no cost.

Consider first the model where $\kappa > 0$ and $q = 0$, so that the individual receives no costless adjustment opportunities. In this case, which I refer to as the (L, c, U) model, the optimal income decisions of an individual are driven by the evolution of his productivity θ and the tax rates that he faces, that is, labor supply considerations. The individual adjusts whenever his desired income y^* gets much larger (he works too little) or much lower (he works too much) than his actual income y . This model is an impulse control problem similar to those analyzed in the operations research or investment literatures, e.g., Harrison, Sellke, and Taylor (1985), Dixit and Pindyck (1994), Stokey (2004). Suppose next that the fixed cost is infinite, $\kappa = \infty$, and the individual receives costless adjustment opportunities at rate $q > 0$. The interpretation of this case is that individuals are permanently searching to adjust their income (at no cost), but can only do so when they exogenously receive an offer. Thus their behavior is entirely dictated by the arrival of exogenous adjustment opportunities, i.e., the demand side of the labor market. This model is similar to that of Calvo (1983) in the monetary literature. In the general model with $\kappa, q \in (0, \infty)$, both margins are active, and the relative size of the two parameters determines the extent to which individual income adjustment decisions are endogenously chosen (labor supply) or exogenously received (labor demand). This model is similar to the CalvoPlus model of Nakamura and Steinsson (2010) and Alvarez, Le Bihan, and Lippi (2014).

Impulse control problem. I now set up the individual’s problem, who must decide when and by how much to adjust his income in response to changes in his productivity or taxes. An impulse control is defined as a sequence of stopping times $0 \equiv t_0 \leq t_1 \leq \dots \leq t_i \leq \dots$ adapted to $\{\mathcal{F}_t\}$, and a sequence of random variables $\Upsilon_1 \leq \dots \leq \Upsilon_i \leq \dots$ measurable with respect to the minimum σ -algebra of events up to t_i , $\{\mathcal{F}_{t_i}\}$. These represent respectively the timing and

⁴Note that the specification of the utility function (2) implies that the hours that the individual spends resting on the job are treated as “leisure”.

the size of the (log-) labor income adjustments, so that $y_t = \Upsilon_{t_i} \times y_{t_{i-1}}$ for all $t \in [t_i, t_{i+1})$. At time 0 an individual with current productivity θ_0 and current income y_0 solves the following sequential maximization problem:

$$\max_{\{t_i, y_{t_i}\}_{i=1}^{\infty}} \mathbb{E}_{\theta_0} \left[\sum_{i=0}^{\infty} \int_{t_i}^{t_{i+1}} e^{-(\rho+\beta+q)t} U(\theta_t, y_{t_i}) dt - \sum_{i=1}^{\infty} e^{-(\rho+\beta+q)t_i} \kappa(\theta_{t_i}) \right], \quad (28)$$

subject to the law of motion (1) of productivity. The arrival rate q of costless adjustment opportunities acts only as an additional discount rate in this problem, because the occurrence of this event triggers an optimal and costless adjustment and thus renders past income decisions irrelevant.

Let $\mathcal{V}(\theta_0, y_0)$ denote the value function of the individual with state (θ_0, y_0) , that is, the expected present discounted value of his lifetime utility net of the adjustment costs, given his optimal income decisions. In the sequel I set up the recursive formulation of the individual's problem (28), and formally define and characterize the value function $\mathcal{V}(\theta_0, y_0)$.

State variables and homogeneity. The impulse control problem that the individual solves is given by (28), and depends on two state variables: the individual's exogenous productivity θ , and the current value of his labor income y . Define the *income deviation* δ as the log-difference between the actual and desired taxable incomes y and y^* respectively, that is

$$\delta_t \equiv \ln(y_t) - \ln(y_t^*). \quad (29)$$

While the individual does not adjust his income, the deviation evolves according to the following process:

$$d\delta_t = -d \ln y_t^* = -\mu_y dt - \sigma_y d\mathcal{W}_t, \quad (30)$$

so that the income deviation δ and the frictionless log-income $\ln y^*$ are perfectly negatively correlated. In the sequel I use (y^*, δ) as the two state variables of the problem, that is, the individual's optimal frictionless income and the deviation of his actual income away from this optimum. Accordingly, the individual's value function is denoted by $\mathcal{V}(y^*, \delta)$.

Replacing θ by its expression (6) as a function of y^* in the definition of the utility function (2), it is easy to show that the flow utility is homogeneous in desired disposable income c^* ,

$$U(y^*, \delta) = \frac{1-\tau}{1-p} (y^*)^{1-p} \left[e^{(1-p)\delta} - \frac{1-p}{1+1/\varepsilon} e^{(1+1/\varepsilon)\delta} \right] \equiv c^* u(\delta). \quad (31)$$

A second-order Taylor approximation of the utility function per unit of consumption $u(\delta)$ around the frictionless optimum $\delta = 0$ implies that the utility loss from failing to optimize is locally quadratic around the frictionless income y^* , with a curvature proportional to $(1-p)(p + \frac{1}{\varepsilon})$. This locally second-order cost of deviating from the frictionless optimum will generate large

inaction regions for even small adjustment costs.

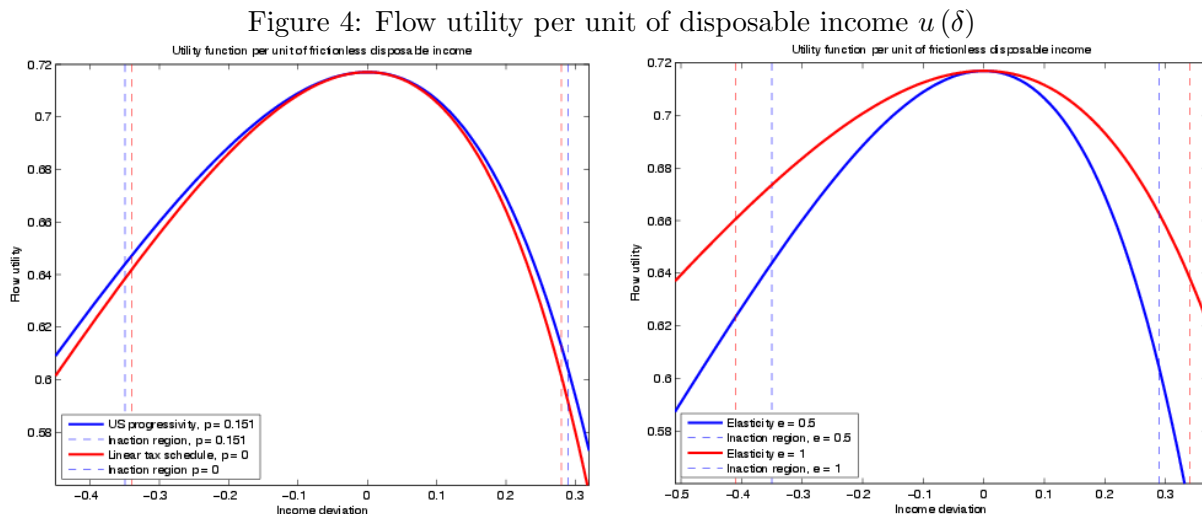
The following lemma allows us to reduce the dimension of the state space.

Lemma 1. *The value function of an individual with desired income y^* and deviation δ is homogeneous of degree one in c^* . That is, it can be written as*

$$\mathcal{V}(y^*, \delta) = \left[\frac{1 - \tau}{1 - p} (y^*)^{1-p} \right] v(\delta) = c^* \times v(\delta). \quad (32)$$

Proof. See Appendix. □

Figure 4 shows the flow utility per unit of frictionless income $u(\delta)$ (along with the optimal inaction region, characterized below). In the first panel, the utility function is plotted for two values of the rate of progressivity p , namely $p = 0.151$ (U.S. tax schedule) and $p = 0$ (linear tax schedule), where the latter is shifted upward so that the maxima of the two curves coincide. In the second panel, the utility function is plotted for two values of the elasticity ε , namely $\varepsilon = 0.5$ and $\varepsilon = 1$. Note in particular that a higher labor income elasticity makes the deviation from the optimal frictionless income less costly (and will therefore induce agents to remain inactive longer). In the extreme case where labor supply is fully inelastic, i.e. $\varepsilon = 0$, agents adjust their income at every instant.



The dynamics of the frictionless and homogeneous state variable c^* has been characterized in Section 2.2. I now analyze the dynamics of the deviation state variable δ , and the value function per unit of frictionless income $v(\delta)$.

Optimal impulse control policy. I now describe the optimal individual adjustment behavior. For any income y , the optimal control policy \mathbf{p}^* is characterized by an interval of inaction $(\underline{\delta}, \bar{\delta})$ and a return point δ^* , with $\underline{\delta} < \delta^* < \bar{\delta}$. The corresponding policy is to exert control if

the state process attempts to exit the open interval $(\underline{\delta}, \bar{\delta})$ and the state is uncontrolled when in $(\underline{\delta}, \bar{\delta})$ unless a costless adjustment opportunity is received. When the state process strikes or is below $\underline{\delta}$ or above $\bar{\delta}$, and when an adjustment opportunity is exogenously received, the control instantaneously moves the state to δ^* , hence the individual adjusts his income from y to $y' = ye^{\delta^* - \delta}$. I characterize the optimal individual policy in two steps. First, I describe the solution to the *fixed boundary problem*, i.e., I characterize the value function of an individual who behaves according to a given (not necessarily optimal) control band policy. Second, I describe the solution to the *free boundary problem*, i.e., I provide sufficient conditions under which this particular control band policy is indeed the optimal impulse control.

For any twice differentiable function $f : \mathbb{R} \rightarrow \mathbb{R}$, define the operator

$$\begin{aligned} \mathcal{L}f(x) = & - \left[\rho + \beta - \mu_c - \frac{1}{2}\sigma_c^2 \right] f(x) + q[f(\delta^*) - f(x)] \\ & - \frac{\mu_c + \sigma_c^2}{1-p} f'(x) + \frac{1}{2} \frac{\sigma_c^2}{(1-p)^2} f''(x), \quad \forall x \in \mathbb{R}. \end{aligned} \quad (33)$$

The following proposition characterizes the solution to the fixed boundary problem defined by any admissible control band policy $\{\underline{\delta}, \delta^*, \bar{\delta}\}$.

Proposition 4. *Suppose that \mathbf{p} is an admissible impulse control characterized by the thresholds and target $\{\underline{\delta}, \delta^*, \bar{\delta}\}$. If $v(\cdot)$ is \mathcal{C}^1 in \mathbb{R} and \mathcal{C}^2 in $\mathbb{R} \setminus \{\underline{\delta}, \bar{\delta}\}$, and if it is the solution to the differential equation problem*

$$\mathcal{L}v(\delta) = -u(\delta), \quad \forall \delta \in (\underline{\delta}, \bar{\delta}), \quad (34)$$

with the value matching conditions

$$\begin{aligned} v(\delta) &= v(\delta^*) - \kappa, \quad \forall \delta \leq \underline{\delta}, \\ v(\delta) &= v(\delta^*) - \kappa, \quad \forall \delta \geq \bar{\delta}, \end{aligned} \quad (35)$$

then $v(\delta)$ is the value function $v_{\mathbf{p}}(\delta)$ associated with the control policy \mathbf{p} .

Proof. See Appendix. □

Equation (34) is the Hamilton-Jacobi-Bellman (HJB) equation of the individual's problem. The interpretation of this equation is as follows. Interpreting the entitlement to the flow of incomes and deviations as an asset, and $\mathcal{V}(y^*, \delta)$ as its value, we can write:

$$(\rho + \beta) \mathcal{V}(y^*, \delta) = U(y^*, \delta) + \frac{\mathbb{E}_t[d\mathcal{V}(y^*, \delta)]}{dt} + q[\mathcal{V}(y^*, \delta^*) - \mathcal{V}(y^*, \delta)].$$

The left hand side gives the normal return per unit time that an individual, using $(\rho + \beta)$ as the discount rate, would require for holding this asset. The first term on the right hand side is the immediate payout or dividend from the asset. The second term is its expected rate of capital

gain or loss. The third term is the change in the value of the asset in case a job opportunity is received, which occurs at rate q per unit time. Thus, the right hand side is the expected total return per unit time from holding the asset. The equality is a no-arbitrage condition, expressing the investor's willingness to hold the asset. Using Itô's formula, we can express the second term on the right hand side as a function of the first and second partial derivatives of the value function \mathcal{V} and the drift and volatility of the income and deviation processes. We finally obtain the HJB equation (34) for $v(\delta)$ using the homogeneity of the value function.

Equations (35) are the value-matching conditions, which state that at the boundaries of the inaction region the individual must be indifferent between adjusting his income (and paying the fixed cost) and not adjusting. These boundary conditions pin down the solution to the HJB differential equation. Note that the solution of the fixed boundary problem (i.e., for given values of $\underline{\delta}, \delta^*, \bar{\delta}$) can be easily characterized explicitly: (34) is a second-order ordinary differential equation with constant coefficients, whose general solution is well known; the two constants it depends on are obtained by inverting the 2×2 linear system of two equations and two unknowns (35).

The following proposition characterizes the solution to the free boundary problem, i.e., it provides sufficient conditions under which a control band policy, characterized by (34,35), is the optimal impulse control.

Proposition 5. *Assume that:*

- (i) *There exist $\{\underline{\delta}, \delta^*, \bar{\delta}\}$ s.t. $v(\cdot)$ solves the DE problem (34,35).*
- (ii) *Smooth-pasting conditions: $v'(\cdot)$ is continuous, i.e.,*

$$\begin{aligned} v'(\underline{\delta}^+) &= v'(\underline{\delta}^-) = 0, \\ v'(\bar{\delta}^+) &= v'(\bar{\delta}^-) = 0. \end{aligned} \tag{36}$$

- (iii) *Optimality condition: δ^* maximizes $v(\cdot)$ on $(\underline{\delta}, \bar{\delta})$, i.e.,*

$$v'(\delta^*) = 0. \tag{37}$$

Then $v(\delta)$ is the value function $v_{\mathbf{p}^}(\delta)$ and $\mathbf{p}^* = \{\underline{\delta}, \delta^*, \bar{\delta}\}$ is the optimal policy.*

Proof. See Appendix. □

Proposition 5 shows that a control band policy, whose value function is characterized by (34,35) in Proposition 4, is the optimal impulse control if its value function satisfies in addition the smooth-pasting conditions (36), which equate the marginal value and the marginal cost of adjusting income, and the optimality condition (37), which sets the optimal return point to the maximum of the value function. The optimal individual adjustment policy is completely characterized by equations (34) to (37).

Note finally that an individual who is born at time 0 with productivity θ_0 , and hence frictionless income $y_0^* \equiv y^*(\theta_0)$, optimally chooses to start working with the deviation δ^* , i.e., with the actual labor income $y_0 = y_0^* e^{\delta^*}$.

Effects of tax policy on individual behavior. The parameter τ has no effect on the optimal adjustment policy $\{\underline{\delta}, \delta^*, \bar{\delta}\}$. An increase in p , however, lowers the volatility σ_y of the income process, which reduces the option value of waiting to adjust, and leads to a *narrower* inaction region. Thus a higher rate of progressivity has an ambiguous effect on the frequency of adjustment T_a : on the one hand, the lower volatility makes individuals reach the boundaries of their inaction region and adjust less often; on the other hand, the inaction region is narrower, which tends to make them adjust more often. The first panel of Figure 5 shows the value function of deviations $v(\delta)$, both in the (L, c, U) limit (see calibration in Section 5.4), and in a CalvoPlus model where the fixed cost is three times larger and the rate of arrival of job offers q is strictly positive. In both cases, the maximum of the value function is at δ^* , the optimal deviation chosen upon adjustment. The value function is hump-shaped within the inaction region: the farther an individual is from his optimal frictionless income (i.e., he works too much if δ is large, or too little if $-\delta$ is large), the worst off he is. Outside the inaction region, the value function is flat: an individual with deviation $\delta < \underline{\delta}$ or $\delta > \bar{\delta}$ adjusts immediately by paying the fixed cost, so is equally well off at any point outside the inaction region. From the value matching conditions, the difference between the maximum (at δ^*) and the minima (at $\underline{\delta}$ and $\bar{\delta}$) of the value function is equal to the value of the value of the fixed cost κ . Finally, the value function and its derivative are continuous at the frontiers of the inaction region (value-matching and smooth-pasting conditions). In the Appendix, I provide an algorithm to compute the optimal control band policy.

4.2 Aggregation

I now analyze how individuals' optimal decisions characterized in Section 4.1 aggregate to form the income and deviation distributions.

Joint distributions of incomes and deviations. Let $f_{y,\delta}^t(\cdot, \cdot)$ denote the joint distribution of (actual) incomes y and deviations δ at time t . The dynamics of this distribution are described by the following Kolmogorov forward (or Fokker-Planck) equation: for all $y \in \mathbb{R}_+^*$ and all $\delta \in (\underline{\delta}, \bar{\delta}) \setminus \{\delta^*\}$,

$$\frac{\partial f_{y,\delta}^t(y, \delta)}{\partial t} = -(\beta + q) f_{y,\delta}^t(y, \delta) + \mu_y \frac{\partial f_{y,\delta}^t(y, \delta)}{\partial \delta} + \frac{\sigma_y^2}{2} \frac{\partial^2 f_{y,\delta}^t(y, \delta)}{\partial \delta^2}. \quad (38)$$

The Kolmogorov forward equation (38) has the following interpretation. At a given income level y , the joint density at the deviation $\delta \in (\underline{\delta}, \bar{\delta}) \setminus \{\delta^*\}$ is reduced by the fraction of individuals with deviation δ who move away from there, and is increased by individuals who start from other

deviations $\delta' \neq \delta$ and end up at δ , following either an increase in their productivity if $\delta' > \delta$ (so that their desired income y^* increases, and they would now like to work more), or a decrease in their productivity if $\delta' < \delta$. These flows occur both because of the general drift μ_y and the volatility σ_y of individual productivity, and are summarized by the last two terms of (38). Moreover, at any income and deviation level (y, δ) , the distribution loses mass at rate β (due to the movements out of the labor force) plus q (due to the exogenous adjustment opportunities). These flows must equal on net the change in the density at (y, δ) , i.e., the left hand side of (38). In particular, in the steady state, these flows in and out of (y, δ) balance exactly, so that the left hand side is equal to zero, and the solution to this partial differential equation is the stationary distribution $f_{y,\delta}(y, \delta)$. Note that this equation does not hold at δ^* where the inflow from non-employment and from endogenous and exogenous adjustments produces a kink in the density.

The Kolmogorov forward equation is subject to the following boundary conditions. For all t , the density function $f_{y,\delta}^t(\cdot, \cdot)$ sums to one. It is continuous on $\mathbb{R}_+^* \times (\underline{\delta}, \bar{\delta})$, so that for all $y \in \mathbb{R}_+^*$, $f_{y,\delta}^t(y, \delta^{*-}) = f_{y,\delta}^t(y, \delta^{*+})$. The boundaries $\underline{\delta}$ and $\bar{\delta}$ are absorbing, so that there is no mass at the edges of the inaction region: for all $y \in \mathbb{R}_+^*$,

$$f_{y,\delta}^t(y, \bar{\delta}) = f_{y,\delta}^t(y, \underline{\delta}) = 0.$$

Intuitively, this is because individuals who reach a boundary of their inaction region adjust immediately within the region. Finally, total flows in and out of the deviation level δ^* must balance. This is described by the following conservation equation, which for simplicity I write for the density of log-incomes $f_{\ln y, \delta}^t(\hat{y}, \delta) = e^{\hat{y}} f_{y, \delta}^t(e^{\hat{y}}, \delta)$: for all $\hat{y} \in \mathbb{R}$,

$$\begin{aligned} & \beta f_{\ln y_0^*}(\hat{y}) + q \int_{\underline{\delta}}^{\bar{\delta}} f_{\ln y, \delta}^t(\hat{y} + \{\delta - \delta^*\}, \delta) d\delta \\ = & \frac{\sigma_y^2}{2} \left[\frac{\partial f_{\ln y, \delta}^t}{\partial \delta}(\hat{y}, \delta^{*-}) - \frac{\partial f_{\ln y, \delta}^t}{\partial \delta}(\hat{y}, \delta^{*+}) \right] \\ & + \frac{\sigma_y^2}{2} \left[\frac{\partial f_{\ln y, \delta}^t}{\partial \delta}(\hat{y} + \{\bar{\delta} - \delta^*\}, \bar{\delta}^-) - \frac{\partial f_{\ln y, \delta}^t}{\partial \delta}(\hat{y} + \{\underline{\delta} - \delta^*\}, \underline{\delta}^+) \right]. \end{aligned}$$

The interpretation of this equation is as follows. The density at δ^* has a kink, because of the inflow of agents at this point coming from adjusters from inside and outside of the labor force. The left hand side is the total inflow at the point (\hat{y}, δ^*) coming from individuals joining the labor force (β) at income level \hat{y} , and from individuals adjusting their income from deviation $\delta \neq \delta^*$ after receiving a costless adjustment opportunity (q). These individuals adjust from income $\hat{y} + \{\delta - \delta^*\}$ and deviation δ to income \hat{y} and deviation δ^* (so that their frictionless log-income \hat{y}^* is equal to $\hat{y} - \delta^*$ both just before and just after the adjustment). The first term on the right hand side describes the flows in and out of the point (\hat{y}, δ^*) coming from individuals

with deviation just below or just above δ^* , due to an increase or decrease in their productivity. The second term on the right hand side describes the flows into δ^* coming from the individuals who adjust because their deviation crosses the thresholds $\bar{\delta}$ and $\underline{\delta}$ of the inaction region. Their log-income thus moves from $\hat{y} + \{\bar{\delta} - \delta^*\}$ and $\hat{y} + \{\underline{\delta} - \delta^*\}$, respectively, to \hat{y} .

In the Appendix I characterize further the stationary joint density of incomes y and deviations δ , $f_{y,\delta}(\cdot, \cdot)$, which is obtained by equating the left-hand side of (38) to zero. In particular, letting $r_{1,\delta} < 0 < r_{2,\delta}$ denote

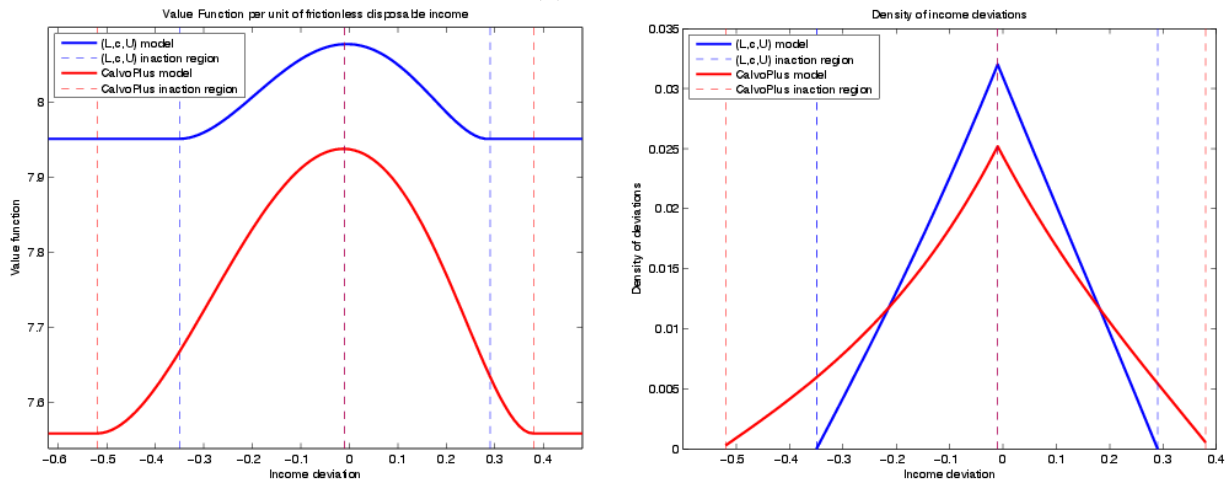
$$r_{1,\delta}, r_{2,\delta} = -\frac{\mu_y}{\sigma_y^2} \pm \sqrt{\left(\frac{\mu_y}{\sigma_y^2}\right)^2 + \frac{2}{\sigma_y^2}(\beta + q)},$$

the steady-state density of deviations conditional on an (actual) income level y is independent of y and is given by

$$f_\delta(\delta) = a_\delta \left[\frac{e^{-r_{1,\delta}(\delta - \underline{\delta})} - e^{-r_{2,\delta}(\delta - \underline{\delta})}}{e^{-r_{1,\delta}(\delta^* - \underline{\delta})} - e^{-r_{2,\delta}(\delta^* - \underline{\delta})}} \mathbb{I}_{[\underline{\delta}, \delta^*]} + \frac{e^{r_{2,\delta}(\bar{\delta} - \delta)} - e^{r_{1,\delta}(\bar{\delta} - \delta)}}{e^{r_{2,\delta}(\bar{\delta} - \delta^*)} - e^{r_{1,\delta}(\bar{\delta} - \delta^*)}} \mathbb{I}_{[\delta^*, \bar{\delta}]} \right], \quad (39)$$

where $a_\delta \in \mathbb{R}$ is a scaling constant which ensures that $f_\delta(\delta)$ sums to one. The second panel of Figure 5 shows the stationary density of deviations $f_\delta(\delta)$ in the (L, c, U) model and in a CalvoPlus model where the fixed cost is three times larger and the rate of arrival of job offers q is strictly positive. It has a kink at δ^* , and the boundaries $\underline{\delta}$ and $\bar{\delta}$ are absorbing.

Figure 5: Value function $v(\delta)$ and stationary density of deviations



5 Effects of taxes in the frictional model

In this section, I study the government's public good provision and redistribution problems described in Section 3, in the frictional economy where the individual and aggregate behaviors are as described in Section 4. In Section 5.1, I define the marginal social welfare weights in the

context of the frictional model, which are useful to describe the welfare effects of tax changes, and show how they differ from the frictionless weights defined in Section 3.2. In Sections 5.3, 5.4, and 5.5, I study and quantify the *long run* effects of tax policy on social welfare in the public good provision and the redistribution problems, and contrast them with the corresponding effects obtained in the frictionless model. Finally, in Section 5.2, I analyze the *transition path* of the economy after a tax change until the new long run steady state is reached.

5.1 Marginal social welfare weights

Government objective. The government chooses the tax schedule $T = \{\tau, p\}$. I consider the same long-run social welfare criteria as in Section 3, \mathcal{W}_{pg} and \mathcal{W}_{r} . Letting $f_{\theta,y}(\cdot)$ (resp., $F_{\theta,y}$) denote the stationary joint density (resp., c.d.f.) of productivities θ and taxable incomes y , the social objective in both equations (16) and (17) is now replaced by

$$\int_0^\infty \int_0^\infty \mathcal{G}(\mathcal{V}(\theta, y)) dF_{\theta,y}(\theta, y).$$

This is because there are now two dimensions of heterogeneity in the population, namely the desired (frictionless) income (y^* , or productivity θ) and the actual income (y , or the income deviation δ). Importantly, the tax schedule is still a tax on income y only; in particular, the definition of tax revenue is identical to (15). I denote by λ the marginal value of public funds in the frictional model.

Individual welfare in the frictional model. In the frictional model, we can use equations (8) and (32) to show that the value function of an individual with *actual* income y and deviation δ is given by

$$\tilde{\mathcal{V}}(y, \delta) \equiv \mathcal{V}(ye^{-\delta}, \delta) = \mathcal{V}^*(y) \times \tilde{v}(\delta), \quad (40)$$

where $\mathcal{V}^*(y)$ is the value function of an individual who would earn income y in the frictionless model, and $\tilde{v}(\delta)$ is the value of having a deviation δ *conditional on an actual income level*, defined by:

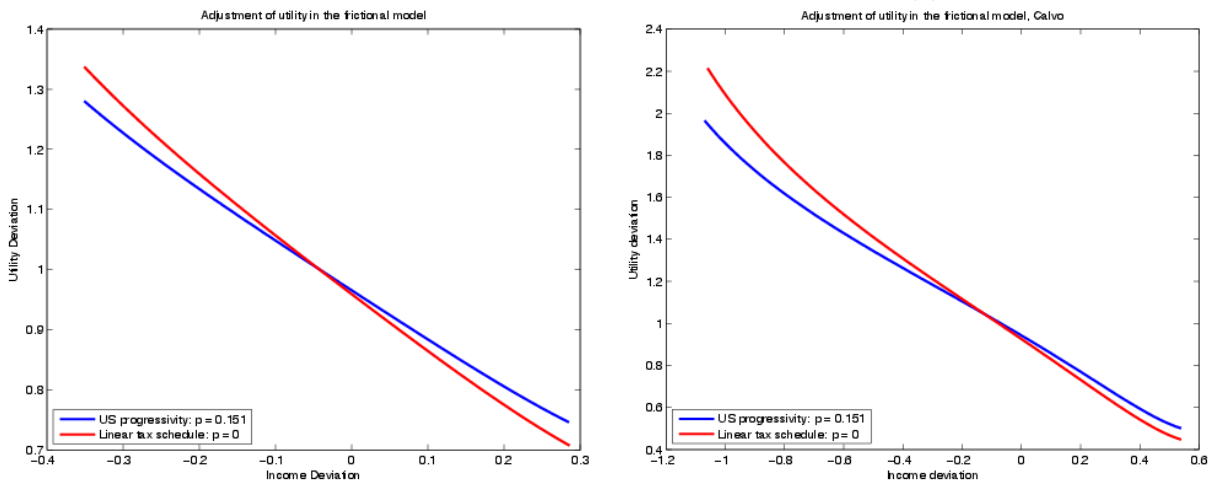
$$\tilde{v}(\delta) = \left(\frac{1}{\rho + \beta - \mu_c - \frac{1}{2}\sigma_c^2} \frac{1 + p\varepsilon}{1 + \varepsilon} \right)^{-1} e^{-(1-p)\delta} v(\delta). \quad (41)$$

Writing the value function as in (40) shows that the indirect lifetime utility of an individual with income y in the frictional model is equal to the utility he would get in the frictionless setting if he earned the same level of income y , times the scaling factor $\tilde{v}(\delta)$ which depends on the current deviation away from his desired income. The value of deviation conditional on income $\tilde{v}(\delta)$ is the solution to an HJB equation derived formally in the Appendix. It is decreasing in the deviation δ , as individuals who earn income y but need to provide less effort have a higher utility than those who earn the same income but need to provide more effort (i.e., those who work too hard to earn this income). This is the main difference between the frictionless and

the frictional models: in the former case, there is a representative agent at each income level y , while in the latter case the population constituting an income group is heterogeneous.

Note that given the optimal policy $\{\underline{\delta}, \delta^*, \bar{\delta}\}$, it is easy to characterize $\tilde{v}(\delta)$ in closed-form as the solution to a fixed boundary problem, as explained in Section 4.1. Importantly, $\tilde{v}(\delta)$ is endogenous to tax policy; specifically, it depends on the progressivity parameter p . Figure 5 plots the value function conditional on income $\tilde{v}(\delta)$. The graph in the first panel is for the (L, c, U) limit (with $q = 0$), while the graph in the second panel is for the model where the fixed adjustment cost is 10 times larger than in the (L, c, U) limit (with $q > 0$, so that the frequency of adjustment is unchanged), i.e., closer to the Calvo limit. The graphs show $\tilde{v}(\delta)$ for two values of p , namely $p = 0.151$ (U.S. tax schedule) and $p = 0$ (linear tax schedule). It shows that for a given deviation, a higher degree of progressivity makes people who put more (resp., less) effort better (resp., worse) off conditional on income. This comes from the fact that the cost of effort, conditional on earning an income c , is proportional to $(1 - p) e^{(p+1/\varepsilon)\delta}$ (see equation (31)). This expression shows that there are two competing effects of an increase in progressivity on the cost of effort for a given deviation, but that on net the cost decreases with progressivity. This property of the utility function translates into the shape of the value function.

Figure 6: Value function conditional on income $\tilde{v}(\delta)$



The frictional model set up in Section 4 thus provides a tractable microfoundation of the *heterogeneous* distribution of individual utilities within each income group y that arises due to the existence of income adjustment frictions, as well as its *endogeneity* to tax policy.

Marginal social welfare weights in the frictional model. I now define the marginal social welfare weights in the frictional model, which are important to characterize the effects of tax changes on social welfare. Keeping in mind the relationship (40) between the individual value functions in the frictionless and the frictional models, consider the same thought experiments as in Section 3.2. That is, take an individual with initial income and deviation y_0 and δ

respectively, and compute the effect on the present discounted value of his utility of giving him an additional consumption stream $\{\hat{c}_t\}_{t \geq 0}$ which evolves according to the same process (μ_c, σ_c) as his frictionless disposable income c_t^* . His value function given this additional consumption stream is equal to:

$$\mathcal{V}(y_0, \delta; \hat{c}) = \mathcal{V}^*(y_0; \hat{c}) \times \tilde{v}(\delta), \quad (42)$$

where $\mathcal{V}^*(y_0; \hat{c})$ is the corresponding utility he would get if he earned the income y_0 in the *frictionless* model, derived in (18). The value of this additional consumption stream in the frictional model is thus weighted by the value of deviation conditional on income, $\tilde{v}(\delta)$. Therefore, the additional units of consumption do not have the same effect on the welfare of individuals who earn the same income y_0 but have different deviations δ (i.e., who have different productivities), the correction term being exactly the function $\tilde{v}(\delta)$.

The long-run marginal social welfare weights at income level y_0 in the frictional model, $\gamma_{\hat{c}}(y_0)$, are defined as social value of distributing the additional stream $\{\hat{c}_t\}_{t \geq 0}$ *uniformly* among all the individuals who earn the same income y_0 . This social value is expressed in monetary units by dividing it by the marginal value of public funds λ , and is computed when the economy has reached its steady state. Therefore, the welfare weight is given by the *average* over deviations δ of the social welfare increases of individuals (y_0, δ) due to the higher consumption. That is,

$$\frac{\gamma_{\hat{c}}(y_0)}{\lambda} \equiv \frac{1}{\lambda} \int_{\underline{\delta}}^{\bar{\delta}} \left[\frac{d\mathcal{V}(y_0, \delta; \hat{c})}{d\hat{c}} \Big|_{\hat{c}_0=0} \mathcal{G}'(\mathcal{V}(y_0, \delta)) \right] f_{\delta}(\delta) d\delta = \frac{1}{\lambda} \frac{\mathbb{E}_{\delta} [\mathcal{G}'(\mathcal{V}(y_0, \delta)) \tilde{v}(\delta)]}{\rho + \beta - \mu_c - \frac{1}{2}\sigma_c^2}, \quad (43)$$

where $f_{\delta}(\delta)$ is the stationary density of income deviations, derived in (39). Note that $f_{\delta}(\delta)$ is independent of y_0 . In this equation, λ is the marginal value of public funds in the frictional model. It is equal to $\mathcal{R}'(\mathcal{R})$ in the public good provision problem. In the redistribution problem of the frictional model, we will see below that λ is given by an expression similar to (27), where the frictional weights $\gamma_{\hat{c}}(y_0)$ replace the frictionless weights $\gamma_{\hat{c}}^*(y_0)$.

Similarly, the effect on the individual (y_0, δ) 's welfare of increasing the growth rate g_c of his future consumption process induces a change in the p.d.v. of his utility given by:

$$\frac{d\mathcal{V}(y_0, \delta)}{dg_c} = \frac{d\mathcal{V}^*(y_0)}{dg_c} \times \tilde{v}(\delta), \quad (44)$$

where $\frac{d\mathcal{V}^*(y_0)}{dg_c}$ is the change in the corresponding value function at income y_0 in the frictionless model, derived in (20). The marginal social welfare weight at income level y_0 associated with an increase in the growth rate of future consumption, $\gamma_{g_c}(y_0)$, is thus defined by taking the average among all the individuals who earn income y_0 of the social welfare gains generated by these higher consumption growth rates:

$$\frac{\gamma_{g_c}(y_0)}{\lambda} \equiv \frac{1}{\lambda} \int_{\underline{\delta}}^{\bar{\delta}} \left[\frac{d\mathcal{V}(y_0, \delta)}{dg_c} \mathcal{G}'(\mathcal{V}(y_0, \delta)) \right] f_{\delta}(\delta) d\delta = \frac{1}{\lambda} \frac{1 + p\varepsilon}{1 + \varepsilon} \frac{\mathbb{E}_{\delta} [\mathcal{G}'(\mathcal{V}(y_0, \delta)) \tilde{v}(\delta)]}{(\rho + \beta - \mu_c - \frac{1}{2}\sigma_c^2)^2} c_0. \quad (45)$$

Suppose in particular that the social welfare function is CRRA, i.e., $\mathcal{G}(u) = \frac{1}{1-\alpha}u^{1-\alpha}$, for some $\alpha > 0$. In this case, the frictionless and frictional marginal social welfare weights, defined respectively in (19) and (43), are related through

$$\gamma_{\hat{c}}(y) = \gamma_{\hat{c}}^*(y) \times \left[\int_{\underline{\delta}}^{\bar{\delta}} \tilde{v}(\delta)^{1-\alpha} f_{\delta}(\delta) d\delta \right]. \quad (46)$$

In general, the term in brackets is different from 1, so that the marginal social welfare weights that I defined are different in the models with and without frictions, i.e., $\gamma_{\hat{c}}(y) \neq \gamma_{\hat{c}}^*(y)$. Moreover, in the redistribution problem, the marginal value of public funds in the frictional model, λ , is not equal to the marginal value of public funds in the frictionless model, λ^* .

To summarize this discussion, the main difference between the frictionless and the frictional models is that in the former environment, each income group y consists of a representative agent, while in the latter case this group consists of a population of heterogeneous individuals, as the least productive of them provide more effort than the others to earn the same income, and hence have a lower utility. Importantly, the heterogeneity in their value function is an equilibrium object and is endogenous to tax policy. The key is to note that the distribution of individuals within a given income group y is non-degenerate *even in the long-run steady-state* when everyone in the economy has had the opportunity to adjust their income to the tax schedule. I now show that these observations have important consequences on the measurement of the welfare effects of raising taxes.

5.2 Marginal tax rates and provision of public goods

In this section, I consider the public good provision problem (16). I show:

Proposition 6. *In the frictional model, the long-run effect of a uniform increase in marginal tax rates on social welfare is given by*

$$\Gamma_{\text{r}} = 1 - \mathbb{E} \left[\frac{\gamma_{\hat{c}}(y)}{\lambda} \frac{\mathbf{t}_{\tau}(y)}{\mathbb{E}[\mathbf{t}_{\tau}(y)]} \right] - \int_0^{\infty} \frac{T'(y)}{1 - T'(y)} \mathbf{e}_{y,1-T'}^* \frac{y \mathbf{t}'_{\tau}(y)}{\mathbb{E}[\mathbf{t}_{\tau}(y)]} f_y(y) dy, \quad (47)$$

where $\gamma_{\hat{c}}(y)$ denotes the marginal social welfare weight at income level y , defined in (43), and where λ is the marginal value of public funds, equal to $\lambda = \mathcal{R}'(\mathcal{R})$.

Proof. See Appendix. □

Formula (47) shows that the welfare effect of an increase in the marginal tax rate τ in the frictional economy is formally identical to the effect of the same perturbation in the frictionless economy, given by equation (23).

The first important result that Proposition 6 shows is that the relevant labor income elasticity that determines the long run behavioral effect of tax policy is the elasticity of desired, or *frictionless*, income y^* , even though there is a non-degenerate distribution of actual incomes y

around y^* in the frictional model, and despite the fact that the individuals' *actual* elasticities are either zero (as long as they don't adjust) or infinite (at the time they adjust). Intuitively, in the long-run, individuals have had time to fully adjust their incomes to the new tax schedule, and the *structural* elasticity parameter ε drives the magnitude of the *aggregate* response to the tax change. This implies that the individual structural parameter ε is still a *sufficient statistic* for the long run distortionary effects of taxes in the case of a uniform increase in the marginal tax rates (i.e., a change in the parameter τ). In other words, the economy behaves in the long-run as if there were a representative agent at each income level.

The second key result of Proposition 6 is that the only difference between the effect of the tax reform in the frictionless and the frictional economies lies in the fact that the marginal social welfare weights used to characterize the welfare loss of raising taxes are the *frictionless weights* $\gamma_{\hat{c}}^*(y)$ (defined in 19) in the former setting, and the *frictional weights* $\gamma_{\hat{c}}(y)$ (defined in (43)) in the latter. The reason and the key friction in my model is that the income tax instruments available to the government are restricted to taxing and redistributing *the same amount* to all the individuals within a given income group y and cannot condition on the unobservable deviation δ . This is why the relevant marginal social welfare weights in the frictional model are the average (over deviations δ) of the individual effects. To the extent that the average weights in the frictional model aren't equal to the frictionless weights, i.e.,

$$\gamma_{\hat{c}}(y_0) \equiv \int_{\underline{\delta}}^{\bar{\delta}} [\gamma_{\hat{c}}(y_0, \delta)] dF_{\delta}(\delta) \neq \gamma_{\hat{c}}^*(y_0),$$

the welfare effects of tax policy will not be correctly calculated by assuming that the economy is frictionless. This is true even in the long-run steady-state when all the individuals have adjusted their income to the new tax system, because at any point in time there is a non-degenerate distribution of agents at each income level y , who differ in their desired income y^* (i.e., their productivity θ or their deviation δ). Since an income tax schedule is unable to redistribute between individuals who earn the same amount y , the welfare effects of uniformly perturbing the tax rate at income y are not captured in a model where this segment of the population is homogeneous, i.e., consists of a representative agent.

Suppose in particular that the social welfare function is CRRA, i.e., $\mathcal{G}(u) = \frac{1}{1-\alpha}u^{1-\alpha}$, for some $\alpha > 0$. Then the ratio of the frictionless to frictional weights $\gamma_{\hat{c}}(y)/\gamma_{\hat{c}}^*(y)$, derived in (46), is constant in y and increasing in α . That is, the more redistributive the government is, the larger the frictional welfare weights $\gamma_{\hat{c}}(y)$ relative to the frictionless weights $\gamma_{\hat{c}}^*(y)$ at every income level y . This is because conditional on an income level y , the welfare of the agents with a larger deviation δ (that is, those who provide more effort to earn this income, i.e., who work too much) is lower than those with a lower deviation δ , as shown in Figure 6. The more concave the social objective, the higher the welfare loss of raising taxes at any income level y , relative to the effect computed in a frictionless model.

Therefore, the marginal social welfare weights in the frictionless model systematically under-

estimate the true welfare weights at every income level if the planner is redistributive enough. Since in the public good provision problem the marginal value of public funds λ is exogenous and identical in the frictionless and the frictional settings, this implies that the welfare cost of increasing the marginal tax rate to finance the public good, measured by (47), is higher in the frictional than in the frictionless model, and the optimal frictional tax rate (24) is lower than its frictionless counterpart.

5.3 Progressivity and redistribution

In this section, I consider the redistribution problem (16). I show:

Proposition 7. *In the frictional model, the effect on social welfare of an increase in the rate of progressivity p is in general not correctly characterized by formula (26) (where the welfare weights $\gamma_{\hat{c}}^*(y)$ and $\gamma_{\rho}^*(y)$ are replaced by $\gamma_{\hat{c}}(y)$ and $\gamma_{\rho}(y)$), unless the following two conditions hold:*

$$\frac{d \ln \{\underline{\delta}, \delta^*, \bar{\delta}\}}{dp} = \frac{d \ln \sigma_y}{dp} = -\frac{\varepsilon}{1 + p\varepsilon}, \quad \text{and } \forall \delta, \tilde{v}^{(p+dp)} \left[\left(1 - \frac{\varepsilon}{1 + p\varepsilon} dp \right) \delta \right] \Big|_{dp \rightarrow 0} = \tilde{v}^{(p)}(\delta).$$

In particular, the labor income elasticity ε and the marginal social welfare weights $\gamma_{\hat{c}}(y)$, $\gamma_{\rho}(y)$ are not sufficient statistics for the revenue and welfare effects of tax changes.

Proof. See Appendix. □

Proposition 7 shows that the effect on social welfare of an increase in progressivity in the frictional model is in general *not* captured by the sufficient statistic formula (26), unless (i) the semi-elasticities of the target and thresholds of the inaction region with respect to p are equal to that of the volatility of the income process; and (ii) the value function of deviation \tilde{v} is unaffected by the change in progressivity, i.e., is scaled one to one with the size of the inaction region.

First, the formula is generally incorrect if the effect of progressivity on the size of the inaction region does not exactly cancel out its effect on the volatility of the process. If this condition is not satisfied, then the individual structural elasticity $\mathbf{e}_{y,1-T'}^*$ is no longer a sufficient statistic, because the increase in progressivity changes the shape of the density of incomes, and thus the mean value of the deviation δ conditional on y^* . In this case, assuming a representative agent at each income level does not adequately capture the effects of taxes. Note that this condition fails to be satisfied as soon as the frequency of adjustment T_a is affected by tax policy. Intuitively, taxes affect not only desired labor supply, but also the optimal adjustment policy, in this case job turnover. On the other hand, the condition holds in two cases: first, if the tax reform increases the marginal tax rates uniformly (change in τ), as we saw in Proposition 6, because τ does not affect the volatility of the process (and thus neither the size of the inaction region); second, in the Calvo limit (with $\kappa = \infty$ and $q > 0$), as then the frequency of adjustment is exogenous to tax policy. But in the general model, the endogenous option value of waiting to adjust implies that

progressivity affects the optimal adjustment policy, which in turn impacts revenue and welfare. If the reduction in volatility dominates the narrowing of the inaction region, the tax schedule should be less progressive than estimated in the frictionless model, as an increase in progressivity induces an equivalent widening of the inaction region around the frictionless income (relative to the case where the band-size effect exactly cancels out the lower volatility), which impacts welfare negatively.

The second reason why formula (26) is incorrect in the frictional environment is that it fails to account for the *endogeneity* of the value of deviations $\tilde{v}(\delta)$ to taxes. Recall that Figure 6 shows that an increase in progressivity reduces the variance of utilities within each income group, so that a higher degree of progressivity makes the least productive people within an income group better off, and the more productive people worse off. This generates a benefit of *higher* progressivity relative to the frictionless sufficient statistic formula. Models that do not explicitly model the heterogeneity within income groups that arises due to adjustment frictions, and in particular its *endogeneity* to taxes, cannot capture its response to tax changes and miscalculate the welfare effects of tax policy. The marginal social welfare weights (43), even corrected to account for the presence of frictions, are thus not a sufficient statistic for the welfare costs of raising taxes when frictions are present.

Finally that even if these two conditions were satisfied, the relevant marginal social welfare weights would be the *frictional* weights $\gamma_{\hat{c}}(y)$ and $\gamma_{\rho}(y)$ rather than the corresponding frictionless weights $\gamma_{\hat{c}}^*(y)$ and $\gamma_{\rho}^*(y)$. Recall that in the redistribution problem and in the frictionless model, the marginal value of public funds λ^* is defined by equation (27). By Proposition 6, it follows that the marginal value of public funds in the frictional model, λ , is equal to that in the frictionless model, replacing the frictionless weights with the frictional weights. Thus the effects on social welfare of increasing progressivity in a budget-neutral way are equal to those in the frictionless model (besides the two effects discussed above) if the *ratio* of the frictional to frictionless weights, $\gamma_{\hat{c}}(y)/\gamma_{\hat{c}}^*(y)$, is constant in income y . This is satisfied if and only if the social welfare function \mathcal{G} is CRRA, i.e., $\mathcal{G}(u) = \frac{1}{1-\alpha}u^{1-\alpha}$ for $\alpha > 0$. For other social welfare functions, the ratio of frictional to frictionless marginal social welfare weights depends on the income level, so that the welfare effects of changing the rate of progressivity would be miscalculated in the frictionless model. This can lead to a lower or higher effect depending on whether $\gamma_{\hat{c}}(y)/\gamma_{\hat{c}}^*(y)$ is decreasing or increasing with income.

To conclude this section, I showed that the sufficient statistic approach to welfare analysis, exposed for instance by Chetty (2009), has limitations when individual behavior is subject to frictions, as the elasticity parameter and the marginal social welfare weights (even the frictional ones) do not correctly account for the effects of tax changes. The large literature in this field of public finance, for instance Saez (2001) in the static model, Golosov, Tsyvinski, and Werquin (2014) in the dynamic model, shows that formulas of the form (26) hold more generally and irrespective of the underlying (frictionless) model. When adjusting labor supply is costly, however, it appears that we still need to solve structural models after all.

5.4 Quantitative analysis

Calibration of the frictional model. I recalibrate the parameters $(m_y, s_y, \mu_y, \sigma_y)$ of the model so that the income distribution in the frictional economy (given the parameters of the U.S. tax schedule, see Section 3.5) has the same mean and variance of log-incomes $\mathbb{E}[\ln y], \mathbb{V}[\ln y]$ and the same Pareto coefficients at the tails $r_{y,1}, r_{y,2}$ as the empirical income distribution in the U.S. Because in this section I compare the long-run effects of taxation in the frictionless and the frictional models, I keep the same value for the structural elasticity ε as in the frictionless model, namely 0.5 in the benchmark calibration.

The additional parameters to calibrate in the frictional model are the fixed adjustment cost κ , and the arrival rate of costless adjustment opportunities q . In the benchmark (L, c, U) model, where there are no such costless opportunities ($q = 0$), the frequency of adjustments T_a (conditional on remaining in the labor force) is given as a function of the parameters of the inaction region $\underline{\delta}, \bar{\delta}, \delta^*$ by:

$$T_a = \frac{\bar{\delta} - \delta}{\mu_y} \left[\frac{\delta^* - \underline{\delta}}{\bar{\delta} - \underline{\delta}} - \frac{e^{2\delta^*\mu_y/\sigma_y^2} - e^{2\underline{\delta}\mu_y/\sigma_y^2}}{e^{2\bar{\delta}\mu_y/\sigma_y^2} - e^{2\underline{\delta}\mu_y/\sigma_y^2}} \right].$$

There is a one-to-one map between the value of the fixed cost κ and the frequency of adjustment; the larger the adjustment cost, the smaller the frequency T_a . I thus choose κ numerically to match a given value of T_a , which I choose to be equal to three years. In the Calvo model (with $\kappa = \infty$ and $q > 0$), the value of q is again pinned down so that T_a is equal to three years. In the general model with $\kappa < \infty$ and $q > 0$, various combinations of (κ, q) imply that the same frequency of adjustments. A larger value of the fixed cost κ requires a larger value of the rate of arrival q to keep T_a constant. The additional degree of liberty allows me to match another moment in the data, namely the variance of log-income changes for individuals who switch jobs, relative to that of those who stay in their job. Empirically, using SIPP data, the standard deviation of monthly log-earnings (resp., of a four-month average of log-earnings) is 1 log point (resp., 0.84 log points) for the movers, and 0.28 log points (resp., 0.2 log points) for the stayers. In my model, by construction the variance of income is fixed conditional on not adjusting (i.e., on keeping the current job), and positive conditional on adjusting when $q > 0$. For my preferred calibration, I choose the pair (κ, q) so that the standard deviation of log-income changes for movers is equal to 0.7 log points, i.e., roughly the difference between that of the movers and the stayers in the data. I also show the results for the (L, c, U) limit ($q = 0$) and the Calvo limit ($\kappa = \infty$). [The numerical simulations will be updated soon to show the preferred calibration.]

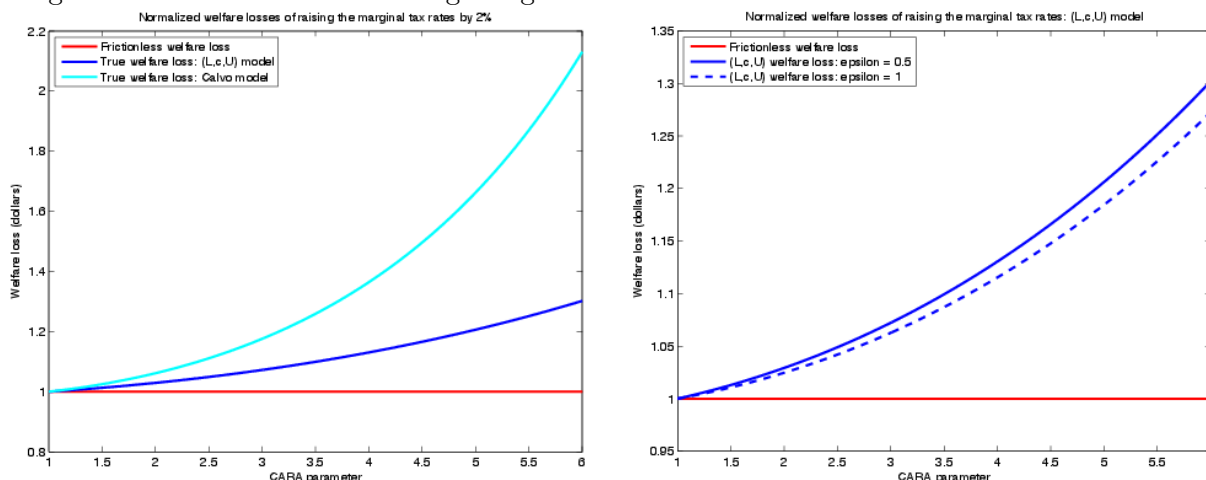
Welfare effects of raising the marginal tax rates in the public good provision problem. Figure 7 shows the social welfare losses of a small increase in marginal tax rates in the frictionless and the frictional models. These are given by formulas (23) and (6), respectively. I compute these effects for a range of values of the CRRA parameter α of the social welfare

function, namely $\alpha \in [1, 6]$. Recall that α summarizes both the individual’s risk aversion coefficient and the concavity of the planner’s social welfare function if he is more redistributive than utilitarian. Rather than taking a stand on the social value of public goods, i.e. the function \mathcal{R} , I normalize the welfare effects for each value of α so that the welfare loss of increasing the tax rates is equal to \$1 in the frictionless model. Using the corresponding value of the marginal value of public funds, I then compute the welfare losses in monetary units associated with the tax increase.

The first panel of Figure 7 compares the welfare losses in the (L, c, U) frictional model (i.e., with $q = 0$) and in the (close to Calvo) model with a fixed cost that is ten times higher (so that $q > 0$), with those in the frictionless model. The simulations show that by ignoring the heterogeneity within each income level, the planner significantly underestimates the costs of raising the marginal tax rates. A \$1 welfare loss estimated using a frictionless view of the economy corresponds to a true welfare loss of \$1.3 in the (L, c, U) model, and of \$2.1 in the Calvo model, when the CRRA parameter of the social welfare function is equal to $\alpha = 6$. The more redistributive the planner, the more the welfare cost of raising taxes is underestimated in the frictionless model.

The second panel compares the welfare losses in the (L, c, U) frictional model with those in the frictionless model, for two values of the labor income elasticity: $\varepsilon \in \{0.5, 1\}$. The frictional losses are closer to the frictionless ones as the elasticity gets larger. This is because a larger elasticity induces individuals to adjust more often their income (see Figure 4). Decreasing the elasticity from $\varepsilon = 1$ to $\varepsilon = 0.5$ when $\alpha = 6$ increases the welfare loss of raising taxes (relative to the frictionless model) by 2.3 percent.

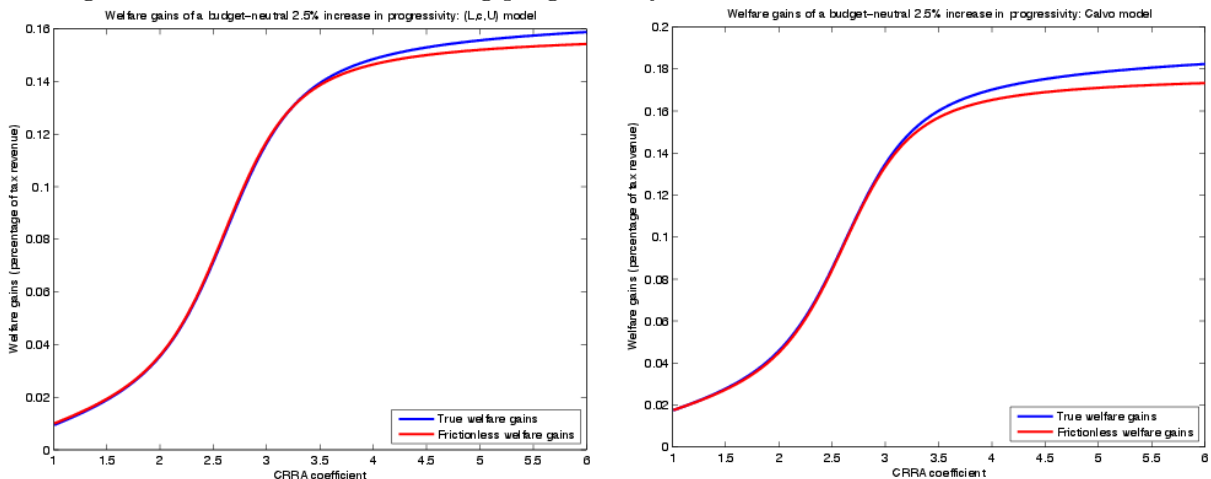
Figure 7: Welfare effects of raising marginal tax rates in the frictional vs. frictionless models



Welfare effects of raising progressivity in the redistribution problem. Figure 8 shows the social welfare effects of a small budget-neutral increase in progressivity in the frictionless and the frictional models, for a CRRA social welfare function. These effects are plotted in

percentage of total government revenue. In the frictionless case, they are given by the sufficient statistic formula (26). If this formula were true in the frictional model (with the modified welfare weights), the estimated gains would coincide, as the ratio of the frictional to the frictionless weights, $\gamma_{\hat{c}}(y) / \gamma_{\hat{c}}^*(y)$, is constant, so that the monetary measure of welfare would be unaffected. However, we saw that this formula ignores (i) the effect of an increase in progressivity on the frequency of adjustment, and (ii) the endogeneity of the distribution of utilities within each income group. Numerically, the former effect is small (because progressivity has only a small effect on the volatility of the process), and the latter dominates: the welfare gains of an increase in progressivity are higher in the *frictional* model if the planner is redistributive enough. In the (L, c, U) benchmark model, the welfare gains of increasing the progressivity of the tax schedule keeping budget constant are 3 percent higher than the planner would estimate by assuming that the economy is frictionless, if $\alpha = 6$. In the Calvo model, these gains are 5.3 percent higher.

Figure 8: Welfare effects of raising progressivity in the frictional vs. frictionless models



5.5 Short-run and long-run labor income elasticities

In this section I analyze the transition path of aggregate income following an unexpected and permanent uniform increase in the marginal tax rates (i.e., an increase in τ), at time 0 when the economy is in its steady state. In the frictionless model, the adjustment of the economy to the new tax schedule would be immediate, as individuals would instantly decrease their labor supply. In the presence of frictions, however, adjustment is sluggish, and an important question is to evaluate the speed of convergence of the economy to the new steady state. The longer it takes, the larger the divergence between the three-year elasticity and the long-run elasticity that matters for the analysis of the long run effects of taxes discussed in Sections 5.1 to 5.3.

The central object to analyze this question is the marginal distribution income deviations $\delta \in [\underline{\delta}, \bar{\delta}]$ within the bands at time t , which I denote by $f_{\delta,t}$. This density satisfies a partial differential (Kolmogorov forward) equation, derived in the Appendix. It is the solution to a

Sturm-Liouville problem, and can be characterized analytically.

Consider first the change in aggregate income on impact, following a change $\Delta\tau$ of the marginal tax rate τ . In response to this perturbation of the tax system, the change in aggregate income is second-order in the tax change, i.e. proportional to $(\Delta\tau)^2$ as $\Delta\tau \rightarrow 0$. This is due to the fact that the frontiers of the inaction region are absorbing (the density shrinks to zero around these boundaries), so that the amount of individuals who adjust their labor supply following the tax shock is small if the shock is small. This is illustrated in the first panel of Figure 9. This implies in particular that the aggregate income elasticity on impact, $\mathbf{e}_{\mathcal{Y},1-\tau}^{(t=0)}$, which is the response of aggregate income to an infinitesimal tax change, is equal to zero. On the other hand, if the tax change is discrete, the elasticity of aggregate income on impact is not equal to zero; it is increasing in the individual structural elasticity $\varepsilon/(1+p\varepsilon)$, increasing in the size of the adjustment $\bar{\Delta}$ (and approximately proportional to $\bar{\Delta}$ for small adjustment costs), and increasing in the size of the tax change $\Delta\tau$.

The following proposition characterizes the impulse response function of aggregate income following the tax change, until convergence to the new steady state.

Proposition 8. *Starting from the steady-state of the economy, the elasticity of aggregate income on impact following a perturbation in the marginal tax rate τ is equal to zero, i.e.,*

$$\mathbf{e}_{\mathcal{Y},1-\tau}^{(t=0)} = 0.$$

At future dates, the law of motion of the aggregate log-income $\hat{\mathcal{Y}} = \mathbb{E}[\ln y]$ is given by:

$$\begin{aligned} \frac{d\hat{\mathcal{Y}}_t}{dt} = & \frac{\sigma_y^2}{2} \{ (\delta^* - \underline{\delta}) [f'_{\delta,t}(\underline{\delta}^+) - f'_{\delta}(\underline{\delta}^+)] + (\bar{\delta} - \delta^*) [f'_{\delta,t}(\bar{\delta}^-) - f'_{\delta}(\bar{\delta}^-)] \} \\ & + (\beta + q) \int_{\underline{\delta}}^{\bar{\delta}} (\delta^* - \delta) [f_{\delta,t}(\delta) - f_{\delta}(\delta)] d\delta. \end{aligned} \quad (48)$$

where $f_{\delta,t}$ denote the density of income deviations δ at time t , and f_{δ} denote the stationary distribution of deviations. In particular, in the Calvo limit, the evolution of the aggregate income elasticity following a perturbation in the marginal tax rate τ is given by:

$$\mathbf{e}_{\mathcal{Y},1-\tau}^{(t)} = \left(1 - e^{-(\beta+q)t}\right) \frac{\varepsilon}{1+p\varepsilon}.$$

Proof. See Appendix. □

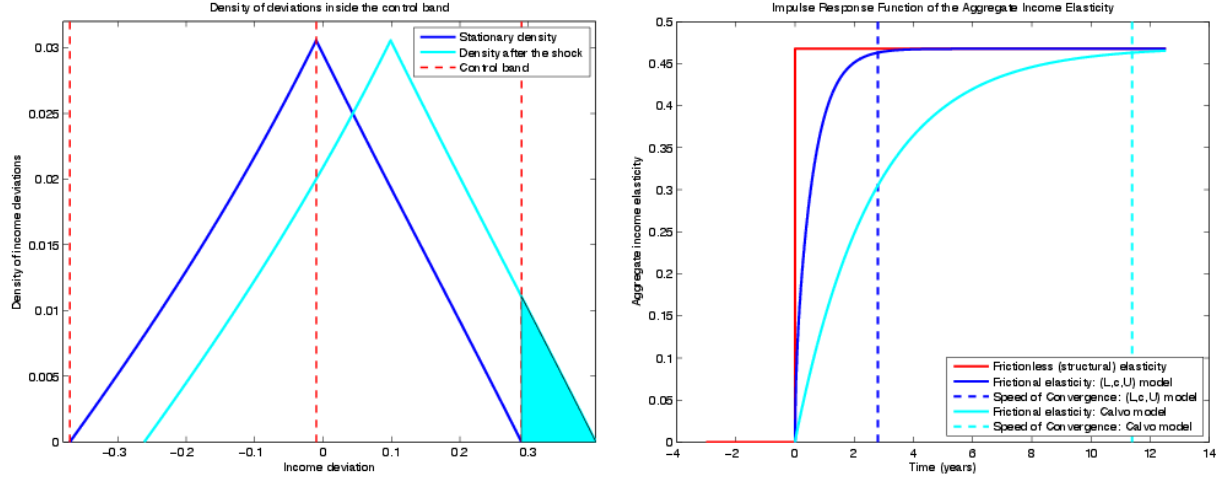
Equation (48) shows that the evolution of aggregate log-income $\hat{\mathcal{Y}}_t$ over time $t > 0$ is determined by the shape of the density of income deviations δ within the inaction region $[\underline{\delta}, \bar{\delta}]$. The first and second terms on the right hand side depend on the density of deviations at the lower boundary $\underline{\delta}$ and at the upper boundary $\bar{\delta}$. Intuitively, the change in aggregate income at a given time t is driven by the fraction of adjusters at this instant, and hence by the density of deviations at the boundaries. Note, however, that the boundaries are absorbing, so that the

density at $\{\underline{\delta}, \bar{\delta}\}$ is always equal to zero. This is why the first-derivative of the density, i.e., the order of the first non-zero term in the Taylor expansion, appears in this equation. Thus, the slopes of the density at the boundaries of the inaction region (relative to the corresponding slopes in the long run stationary distribution), weighted by the size of the upward and downward adjustments, determine by how much aggregate income adjusts at any given point in time in the transition path. The second term on the right hand side of (48) corresponds to the adjustments of individuals from within the inaction region, which are due to the fraction of agents leaving and entering the labor force at each instant (β), and those who receive an exogenous costless adjustment opportunity (q). The change in aggregate income coming from such individuals with deviation δ is given by the density of deviations at this point (relative to the long run density), times the size of the adjustment. Note finally that (48) shows that the path of the aggregate income is differentiable (it has finite variation), even though the sources of variation at the micro level are infinite variation Brownian motions.

For a given adjustment frequency T_a , the (L, c, U) limit (with $\kappa > 0$ and $q = 0$) provides an *upper bound* for the speed of adjustment of aggregate income in response to the tax shock. This is because of a selection effect à la Golosov and Lucas (2007): the individuals who adjust are those whose income is the farthest from their frictionless optimum. They adjust by a large amount, which makes the economy converge fast to its stationary state. At the other extreme, the Calvo limit (with $\kappa = \infty$ and $q > 0$) provides a *lower bound* for the speed of adjustment: in this case, the individuals who adjust at a given instant are a non-selected sample of the population. The fact that there is always a fraction of individuals who are very far from their frictionless optimum makes the adjustment much slower. In the general case, the speed of adjustment is between that of the (L, c, U) and the Calvo limits.

The second panel of Figure 9 plots the evolution of the aggregate income elasticity after the tax shock at time 0. In the frictionless case, the adjustment is immediate, since all individuals adjust instantly to the new tax rates. In the (L, c, U) case, the figure shows that after three years, aggregate income has almost completely adjusted to the tax shock. The three-year elasticity typically estimated in the data is then a good estimate of the long-run elasticity that is relevant for the theory of taxation. In the Calvo case, the aggregate income elasticity takes longer to converge, so that the three-year elasticity is a substantial underestimate (less than two thirds) of the long-run elasticity. In the general model, the more adjustments are driven by labor demand rather than labor supply, i.e. outside the control of the individual rather than optimally chosen, the more the long-run elasticity differs from the short run elasticity.

Figure 9: Density of deviations after the shock and impulse response function of aggregate income



6 Conclusion

In this paper I have set up a novel dynamic framework to analyze the effects of taxes on social welfare when individual labor supply is subject to adjustment frictions. The frictionless model is highly tractable. It can be solved entirely in closed form and is able to match some key features of the empirical income distributions. In this model, the effects of taxes on individual behavior, aggregate distributions, tax revenue, and social welfare are extremely transparent. Most importantly, I use this framework to introduce a fixed cost of adjusting labor supply. I show that the presence of adjustment frictions affects in important ways the measure of the social welfare effects of raising taxes. By ignoring the heterogeneity of utilities within each income group, as the least productive workers have to provide more effort than the others, the frictionless model systematically underestimates the welfare costs of increasing marginal tax rates to redistribute to a public good. By ignoring the endogeneity of the distribution of utilities to taxes, as well as the effects of taxes on the optimal individual adjustment policy, the frictionless model miscalculates the welfare effects of raising the progressivity of the tax schedule, and the standard frictionless sufficient statistic formulas cease to hold.

What are the policy implications of these findings? The key friction in my model is that the available tax instruments are taxes on labor income only, and cannot disentangle more productive from less productive individuals who earn the same income level. Even though the underlying productivities or efforts are unobservable to the planner, the changes in income are observable. The planner can observe the direction, the size and the timing of income adjustments, which reveal whether the individual had been working too hard or too little. Thus, the government should tax not only *income*, but also *changes in income*. If the drift of productivity is positive, so that individuals adjust their income upward on average, there will be more people who are working too little than are working too much at every level of income, and thus taxes should be

higher for those who have earned this income for a while than for those who have just adjusted to this income. I leave it for future research to explore the shape of the optimal tax system when those tax instruments are available.

Another extension that I leave for future research is to set up and estimate a sophisticated structural model of individual behavior, in which productivity and income shocks have not only a permanent but also transitory components, wages are determined endogenously in general equilibrium, the adjustment costs vary with the income level, and the tax and transfer system incorporates more specific elements of the actual tax codes. In such a model, it would be valuable to quantify the theoretical forces highlighted in this paper.

References

- [1] **Altonji, Joseph and Christina Paxson**, “Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-Offs”, *Journal of Labor Economics*, 1988
- [2] **Altonji, Joseph and Christina Paxson**, “Labor Supply, Hours Constraints, and Job Mobility”, *Journal of Human Resources*, 1992
- [3] **Alvarez, Fernando, Hervé Le Bihan, and Franco Lippi**, “Small and Large Price Changes and the Propagation of Monetary Shocks”, *Working Paper*, 2014
- [4] **Alvarez, Fernando and Francesco Lippi**, “Price Setting with Menu Costs for Multi-Product Firms”, *Econometrica*, 2014
- [5] **Alvarez, Fernando and Nancy Stokey**, “Dynamic Programming with Homogeneous Functions”, *Journal of Economic Theory*, 1998
- [6] **Benabou, Roland**, “Tax and Education Policy in a Heterogeneous-Agent Economy: What levels of Redistribution Maximize Growth and Efficiency?”, *Econometrica*, 2002
- [7] **Bertola, Giuseppe and Ricardo Caballero**, “Kinked Adjustment Costs and Aggregate Dynamics”, *NBER Macroeconomics Annual*, 1990
- [8] **Caballero, Ricardo and Eduardo Engel**, “Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach”, *Econometrica*, 1999
- [9] **Calvo, Guillermo**, “Staggered Prices in a Utility-Maximizing Framework”, *Journal of Monetary Economics*, 1983
- [10] **Chetty, Raj**, “Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods”, *Annual Review of Economics*, 2009
- [11] **Chetty, Raj**, “Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply”, *Econometrica*, 2012

- [12] **Chetty, Raj, John Friedman, Tore Olsen, and Luigi Pistaferri**, “Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records”, *Quarterly Journal of Economics*, 2011
- [13] **Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber**, “Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities”, *NBER Macroeconomics Annual*, 2012
- [14] **Chetty, Raj, Adam Looney, and Kory Kroft**, “Salience and Taxation: Theory and Evidence”, *American Economic Review*, 2009
- [15] **Choné, Philippe and Guy Laroque**, “Optimal Taxation in the Extensive Model”, *Journal of Economic Theory*, 2011
- [16] **Cogan, John**, “Fixed Costs and Labor Supply”, *Econometrica*, 1981
- [17] **Diamond, Peter**, “Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates”, *American Economic Review*, 1998
- [18] **Dickens, William and Shelly Lundberg**, “Hours Restrictions and Labor Supply”, *International Economic Review*, 1993
- [19] **Dixit, Avinash and Robert Pindyck**, “Investment under Uncertainty”, *Princeton University Press*, 1994
- [20] **Feng, Haolin and Kumar Muthuraman**, “A Computational Method for Stochastic Impulse Control Problems”, *Mathematics of Operations Research*, 2010
- [21] **Gabaix, Xavier**, “Power Laws in Economics and Finance”, *Annual Review of Economics*, 2009
- [22] **Gelber, Alexander, Damon Jones, and Daniel Sacks**, “Earnings Adjustment Frictions: Evidence from the Social Security Earnings Test”, *Working Paper*, 2013
- [23] **Golosov, Mikhail, and Robert Lucas**, “Menu Costs and Phillips Curves”, *Journal of Political Economy*, 2007
- [24] **Golosov, Mikhail, Aleh Tsyvinski, and Nicolas Werquin**, “A Variational Approach to the Analysis of Tax Systems”, *Working Paper*, 2014
- [25] **Greenwood, Jeremy, Zvi Hercowitz, and Gregory Huffman**, “Investment, Capacity Utilization, and the Real Business Cycle”, *American Economic Review*, 1988
- [26] **Grossman, Sanford and Guy Laroque**, “Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods”, *Econometrica*, 1990

- [27] **Gruber, Jonathan and Emmanuel Saez**, “The Elasticity of Taxable Income: Evidence and Implications”, *Journal of Public Economics*, 2002
- [28] **Guvenen, Fatih, Serdar Ozkan, Fatih Karahan, and Jae Song**, “What Do Data on Millions of US Workers Say About Lifecycle Earnings Risk?”, *Working Paper*, 2014
- [29] **Holmlund, Bertil and Martin Söderström**, “Estimating Dynamic Income Responses to Tax Reforms: Swedish Evidence”, *Working Paper*, 2008
- [30] **Jacquet, Laurence, Etienne Lehmann, and Bruno Van der Linden**, “The Optimal Marginal Tax Rates with both Extensive and Intensive Responses”, *Journal of Economic Theory*, 2013
- [31] **Jones, Charles and Jihee Kim**, “A Schumpeterian Model of Top Income Inequality”, *Working Paper*, 2014
- [32] **Harrison, J. Michael, Thomas Sellke, and Allison Taylor**, “Impulse Control of Brownian Motion”, *Mathematics of Operations Research*, 1985
- [33] **Heathcote, Jonathan, Kjetil Storesletten, and Gianluca Violante**, “Optimal Tax Progressivity: An Analytical Framework”, *Working Paper*, 2014
- [34] **Keane, Michael, and Richard Rogerson**, “Micro and Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom.” *Journal of Economic Literature*, 2012
- [35] **Ljungqvist, Lars and Thomas Sargent**, “A Labor Supply Elasticity Accord?”, *American Economic Review: Papers and Proceedings*, 2011
- [36] **Mirrlees, James**, “An Exploration in the Theory of Optimal Income Taxation”, *Review of Economic Studies*, 1971
- [37] **Meghir, Costas, and Luigi Pistaferri**, “Income Variance Dynamics and Heterogeneity”, *Econometrica*, 2004
- [38] **Meghir, Costas, and Luigi Pistaferri**, “Earnings, Consumption and Life Cycle Choices”, , 2011
- [39] **Musgrave, Richard and Tun Thin**, “Income Tax Progression: 1929-48”, *Journal of Political Economy*, 1948
- [40] **Nakamura, Emi and Jón Steinsson**, “Monetary Non-Neutrality in a Multisector Menu Cost Model”, *Quarterly Journal of Economics*, 2008
- [41] **Nirei, Makoto, and Wataru Souma**, “A Two Factor Model of Income Distribution Dynamics”, *Review of Income and Wealth*, 2007

- [42] **Piketty, Thomas and Emmanuel Saez**, “A Theory of Optimal Inheritance Taxation”, *Econometrica*, 2013
- [43] **Reed, William**, “The Pareto Law of Incomes – An Explanation and an Extension”, *Physica A*, 2003
- [44] **Reed, William and Murray Jorgensen**, “The Double Pareto-Lognormal Distribution – A New Parametric Model for Size Distributions”, *Communications in Statistics - Theory and Methods*, 2003
- [45] **Rogerson, Richard and Johanna Wallenius**, “Micro and Macro Elasticities in a Life Cycle Model with Taxes”, *Journal of Economic Theory*, 2009
- [46] **Rogerson, Richard and Johanna Wallenius**, “Nonconvexities, Retirement, and the Elasticity of Labor Supply”, *American Economic Review*, 2013
- [47] **Saez, Emmanuel**, “Using Elasticities to Derive Optimal Income Tax Rates”, *Review of Economic Studies*, 2001
- [48] **Saez, Emmanuel**, “Optimal Income Transfer Programs: Intensive Versus Extensive Labor Supply Responses”, *Quarterly Journal of Economics*, 2002
- [49] **Saez, Emmanuel, Joel Slemrod and Seth Giertz**, “The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review”, *Journal of Economic Literature*, 2012
- [50] **Saez, Emmanuel and Stefanie Stantcheva**, “Generalized Social Welfare Weights for Optimal Tax Theory”, *Working Paper*, 2014
- [51] **Shourideh, Ali, and Maxim Troshkin**, “Providing Efficient Incentives to Work: Retirement Ages and the Pension System”, *Working Paper*, 2014
- [52] **Stokey, Nancy**, “The Economics of Inaction”, *Princeton University Press*, 2008
- [53] **Toda, Alexis Akira**, “The Double Power Law in Income Distribution: Explanations and Evidence”, *Journal of Economic Behavior and Organization*, 2012

Appendix