Optimal Tax Progressivity in Unionised Labour Markets: What are the Driving Forces?

Stefan Boeters

ZEW

Zentrum für Europäische
Wirtschaftsforschung GmbH

Centre for European
Economic Research

Optimal Tax Progressivity in Unionised Labour Markets: What are the Driving Forces?

Stefan Boeters

Non-technical summary

In the economic policy debate, income tax progressivity is mostly seen as a means of redistribution. The more progressive the tax, the more redistribution from the rich to the poor. However, high tax progressivity also means high marginal tax rates for those with a high income, which leads to labour supply distortions in this group and, as a consequence, to a lower overall potential for redistribution. Under non-competitive wage formation, there are positive aspects of tax progressivity in the efficiency dimension as well, because tax progressivity lowers the incentives for high wage claims and leads to a downward pressure on non-competitive wages. This counteracts the labour supply distortions and creates the potential for a free lunch: simultaneous equity and efficiency gains through higher tax progressivity. I explore the empirical potential for this constellation using a model that includes both wage bargaining and flexible labour supply, so that both sides of the trade-off are captured. It is calibrated to a set of macroeconomic and institutional parameters of large OECD countries, and the optimal degree of tax progressivity is numerically determined.

The most remarkable simulation result is that both at the level of average OECD parameters and for most of the individual countries, optimal tax progressivity is considerably lower than actual progressivity. At the country level, however, we do not get uniform picture. In a few countries optimal progressivity is higher than actual progressivity. The between-country differences can be traced back to differences in the initial conditions. The effect of the initial unemployment rate is particularly strong. The higher initial unemployment, the higher optimal tax progressivity. Another important driver is the general tax level. High taxes in the initial situation lead to a lower optimal level of tax progressivity. The initial level of tax progressivity plays a significant role as well. It affects optimal progressivity through the interaction with labour supply elasticities.

Nicht-technische Zusammenfassung

In der wirtschaftspolitischen Diskussion wird die Progressivität der Einkommensteuer vor allem als ein Instrument der Umverteilung gesehen. Je höher die Progression, desto mehr Umverteilung von Reich zu Arm. Ein hoher Grad der Steuerprogression bedeutet jedoch auch eine hohe marginale Belastung für Bezieher hoher Einkommen. Dies führt zu Verzerrungen des Arbeitsangebots in dieser Gruppe, und in der Konsequenz zu einem geringeren Umverteilungsvolumen. Bei nicht-wettbewerblicher Lohnbildung kommt noch ein weiterer Effekt ins Spiel. Eine höhere Steuerprogression hat dann auch in der Effizienzdimension positive Wirkungen, da sie dämpfend auf die Lohnforderungen wirkt und zu niedrigeren Lohnabschlüssen führt. Dies wirkt den Verzerrungen des Arbeitsangebots entgegen und schafft das Potential eines gleichzeitigen Gewinns in sowohl der Effizienz- als auch der Verteilungsdimension. In diesem Papier untersuche ich mit Hilfe eines numerischen Modells, ob dieses Potential für tatsächliche Volkswirtschaften relevant ist. Das Modell beinhaltet kollektive Lohnverhandlungen und flexibles Arbeitsangebot, so dass beide abzuwägende Seiten erfasst sind. Es wird mit Hilfe makroökonomischer und institutioneller Parameter großer OECD-Volkswirtschaften kalibriert, so dass der optimale Grad der Steuerprogression numerisch bestimmt werden kann.

Das wichtigste Ergebnis der Simulationsrechnungen ist, dass sowohl im OECD-Durchschnitt als auch für die meisten individuellen Länder die optimale Progression deutlich unter der tatsächlichen liegt. Auf der Länder-Ebene ist das Bild allerdings nicht einheitlich: In einigen Ländern liegt die optimale auch über der tatsächlichen Progression. Die Unterschiede zwischen den Ländern können auf Unterschiede in den Ausgangsbedingungen zurückgeführt werden. Dabei ist der Effekt der bestehenden Arbeitslosigkeit besonder stark. Je höher die Arbeitslosigkeit, desto höher die optimale Steuerprogression. Ein zweiter wichtiger Einflussfaktor ist das allgemeine Steuerniveau. Ein hohes Steuerniveau führt zu einer niedrigen optimalen Progression. Schließlich spielt auch das Ausgangsniveau der Steuerprogression eine Rolle. Es beeinflusst die optimale Progression durch die Interaktion mit den Arbeitsangebotselastizitäten in der Kalibrierung des Modells.

Optimal Tax Progressivity in Unionised Labour Markets: What are the Driving Forces?

Stefan Boeters[*]

CPB, Netherlands Bureau for Economic Policy Analysis, Den Haag

October 2009

Abstract

In labour markets with collective wage bargaining higher progressivity of the labour income tax creates a trade-off. On the one hand, wages are lowered and unemployment decreases, on the other hand, the individual labour supply decision is distorted at the hours-of-work margin. The optimal level of tax progressivity within this trade-off is determined using a numerical general equilibrium model with imperfect competition on the goods market, collective wage bargaining and a labour-supply module calibrated to empirically plausible elasticity values. The model is calibrated to macroeconomic and institutional parameters of both the OECD average and a number of individual OECD countries. In most cases the optimal degree of tax progressivity is below the actual level. A decomposition approach shows that the optimal level is increased by high unemployment and by the general tax level.

Keywords: labour taxation, tax progressivity, optimal taxation, collective wage bargaining, unemployment

JEL Code: H21, J22, J51, J64,

[*]Stefan Boeters, CPB, P.O. Box 80510, NL-2508 GM Den Haag, e-mail: s.boeters@cpb.nl. I thank Rob Euwals, Nicole Gürtzgen, Albert van der Horst, Bas Jacobs, Egbert Jongen, Stefanie Schubert and Peter Sørensen for helpful comments on earlier versions of this paper.

1 Introduction

Income tax progressivity creates an efficiency trade-off in labour markets with collective wage bargaining. On the one hand, higher tax progressivity can reduce unemployment, which leads to a welfare gain. On the other hand, the individual labour supply decision is distorted, with an associated deadweight loss. In this paper, I calculate the optimal degree of tax progressivity that results from this trade-off, and I try to identify the driving forces for differences between OECD countries.

The focus on efficiency issues in the discussion of the income tax needs explanation. In the economic policy debate, income tax progressivity is mostly seen as a means of redistribution from the rich to the poor. Traditionally, economist have stressed that efficiency must also be considered in this context. High tax progressivity leads to labour supply distortions in the high-income group and, as a consequence, to a shrinking overall potential for redistribution. The derivation of criteria for an optimum in this situation has been one of the early highlights of optimal taxation theory (Mirrlees, 1971; Tuomala, 1990).

Since the 1980s, extensive research into wage forming mechanisms that are not fully competitive has shown that the equity-efficiency trade-off is not clear-cut. There are positive aspects of tax progressivity in the efficiency dimension as well, because tax progressivity lowers the incentives for high wage claims and leads to a downward pressure on non-competitive wages. This counteracts the labour supply distortions and creates the potential of a free lunch: simultaneous equity and efficiency gains through higher tax progressivity. However, this is only a theoretical potential – whether it is relevant to real-world economies remains a question for empirical research.

The trade-off within the efficiency dimension has often be addressed using an analytical approach, but there are hardly any attempts to quantify it. What is the optimal level of tax progressivity that balances the positive effect of lower unemployment through lower wages with the negative effect on labour supply? This is the question I investigate in this paper using a calibrated, numerical general equilibrium model. I adopt the model of collective wage bargaining between a trade union and an employers' association. This mechanism is dominant in a number of

European countries and has most often been chosen to demonstrate the ambiguous effects of tax progressivity theoretically (Hersoug, 1984; Lockwood and Manning, 1993; Holmlund and Kolm, 1995; Koskela and Vilmunen, 1996).

The numerical model includes both wage bargaining and flexible labour supply, so that both sides of the trade-off are captured. It is calibrated to a set of macroeconomic and institutional parameters of large OECD countries and solved for the optimal degree of tax progressivity. This approach is close in spirit to Sørensen (1999), which is, to my knowledge, the only paper in the literature that comes up with a numerically determined optimal degree of tax progressivity. However, Sørensen's main focus is on the comparison of different mechanisms of wage determination. To provide a sharper focus, he chooses a relatively simple calibration of labour supply, and he fixes the institutional and macroeconomic parameters at stylised values of a typical Western industrialised country. This is where the present paper comes in. The calibration of labour supply is elaborated, so that it accounts for different empirical indicators: labour supply elasticities of hours of work and participation, and with respect to wages as well as to non-wage income.

Why is there only one single paper that tries to quantify the optimal tax progressivity in a labour market with collective wage bargaining? This is most probably explained by the fact that to do so, we are forced to leave the area of general and clear-cut analytical results. No-one has so far come up with illuminating analytical expressions that characterise the optimal point. For an optimal tax analysis that involves two tax rates (in our case: the marginal and the average tax on labour income), we need two indicators per tax: its marginal effect on utility, and its marginal effect on the public budget. The latter soon becomes involved once we include the indirect effects through the changes in the tax bases of other taxes (which is necessary in general equilibrium). It remains possible to derive analytical expressions for these effects, but they no longer provide an insight in the economic mechanisms. Hence the shift to numerical models. Here we lose generality, but we may focus directly on parameters that are *quantitatively relevant* in the situation at hand. Nevertheless, the choice in this paper is to limit the analysis to a *simple* numerical model. The reason is that once we have identified a parameter that is quantitatively important, we do not want to stop at this point, but explain *why* it is important, and why the effect was qualitatively to be expected, even if we could not foresee that it would quantitatively drive the results.

2

In my attempt to exploit the quantitative potential of the model presented in this paper, I calibrate it both to unweighted averages of the institutional and macroeconomic parameters of eight large OECD economies and to the individual country constellations. A decomposition exercise is executed by varying one of the parameters at a time. This allows us to identify the key drivers of the differences in optimal tax progressivity.

The most remarkable simulation result is that both at the level of average OECD parameters and for most of the individual countries, optimal tax progressivity is considerably lower than actual progressivity. At the country level, however, we do not get a uniform picture. In a few countries optimal progressivity is even higher than actual progressivity. The between-country differences can be traced back to differences in the initial conditions. The effect of the initial unemployment rate is particularly strong. The higher initial unemployment, the higher optimal tax progressivity. Another important driver is the general tax level. High taxes in the initial situation lead to a lower optimal level of tax progressivity. The initial level of tax progressivity plays a significant role as well. It affects the optimal progressivity through the interaction with labour supply elasticities. This effect is discussed in detail in the body of the paper.

The model of this paper is set up to focus on one particular trade-off connected with tax progressivity, at the cost of a number of other aspects that are not included. These should be kept in mind when interpreting the results. First, the quantitative results of the paper do not automatically carry over to other theories of unemployment (search-and-matching and efficiency wage theories). However, Pissarides (1998) and Sørensen (1999) show that these approaches produce results similar to the collective bargaining model when they are calibrated to plausible parameter values and applied to taxation issues. The focus on the collective bargaining model is therefore not overly restrictive.

Second, there are other distortions, apart from the effect on labour supply, that run counter to the wage moderating effect of tax progressivity. Examples are Fuest and Huber (1998), who focus on the distortionary effect on human capital formation, Kleven and Sørensen (2004), who describe the effects on dual labour markets, where only one sector is characterised by imperfections, and Koskela and Schöb (2007), who stress the negative effect on workers' effort.

Third, the paper focuses on efficiency issues and abstracts from one core aspect in the early literature on optimal taxation (Mirrlees, 1971; Tuomala, 1990): non-observable productivity differences of heterogeneous agents. This can be seen as an analogue to the Ramsey (1927) approach to indirect taxation, where distributional concerns are ignored as well, in order to get a clear picture of the efficiency dimension.

Finally, the model of this paper remains at the aggregate level and captures labour market institutions and the tax and transfer system only by a small set of macro indicators. Micro econometricians argue that such models miss the very essence of the labour market: heterogeneity. In fact, there are examples of models that combine microeconomically founded mechanisms of involuntary unemployment and demographic as well as institutional heterogeneity in the labour market: Sørensen (1997), Graafland et al. (2001), Aaberge et al. (2004), Arntz et al. (2008). Due to their complexity, the outcomes of such models are often difficult to explain and to decompose into effects that are qualitatively known from the theoretical literature. This interpretation work is simplified through a condensed and simplified "model of the model" (e.g. "Mini-MIMIC" (Bovenberg et al., 2000) as a complement to Graafland et al., 2001). It is in this tradition that the present paper is most appropriately placed.

The plan for the rest of the paper is as follows. In Section 2, I present the different parts of the model, my approach to labour supply calibration, the welfare criterion and the OECD parameters used for the simulations. Sections 3 covers the simulations for the OECD average, different country specifications and systematic parameter variations that make a decomposition of the tax progressivity effect possible. In Section 4, I perform a sensitivity analysis, before Section 5 concludes. The appendix contains the algebraic details of the labour supply calibration and a list of data sources.

2   The model

We consider a small, representative production sector with monopolistic competition in a closed economy. The wage is determined through collective wage bargaining, which produces involuntary unemployment. The government collects taxes on

consumption, profit and labour income. In this situation, the progressivity of the labour income tax is chosen so as to maximise the expected utility of a representative worker.

2.1 Consumption

There are four households, indexed by h: the aggregate households of employed (e) and involuntarily unemployed (u) workers,[1] the recipients of residual income ("capitalists", k) and the public household (g). Consumption of these households, C_h, is composed of a large, but fixed number of symmetric varieties of goods, corresponding to the number of firms in production, n. The varieties are indexed by i with prices p_i, quantities x_i and a uniform consumption tax, t_C. The elasticity of substitution between the individual goods varieties is η.

$$C_h = n^{\frac{1}{1-\eta}} \left(\sum_{i=1}^{n} (x_i^h)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}.$$

This gives demand for individual varieties

$$x_i^h = \frac{C_h}{n} \left(\frac{p_C}{p_i(1+t_C)} \right)^{\eta}. \tag{1}$$

In (1), p_C is the consumer price index

$$p_C = \left(\frac{1}{n} \sum_{i=1}^{n} [p_i(1+t_C)]^{1-\eta} \right)^{\frac{1}{1-\eta}}.$$

so that

$$C_h = \frac{Y_h}{p_C},$$

where

$$Y_e = w(1 - t_L^a)H$$

$$Y_u = cw(1 - t_L^a)H$$

$$Y_k = (1 - t_\pi)\Pi$$

---

[1]These are aggregate households because they represent individual households with varying fixed cost of taking up work (see Section 2.4).

and Y_g is defined in (10). w is the before-tax hourly wage rate, H hours of work, c the replacement rate (see Section 2.3), Π residual income ("profits", $\Pi = \sum \pi_i$), and t_L^a and t_π the average tax rates on labour income and profits, respectively. Adding up the structurally identical demand from all four households, firm i faces the output demand function

$$x_i = \frac{C}{n} \left(\frac{p_C}{p_i(1 + t_C)} \right)^\eta,$$

where

$$C = C_e + C_u + C_k + C_g.$$

2.2   Hours of work

Utility of worker households is assumed to be of the CES type, with consumption, C, and leisure, F, as arguments and elasticity of substitution σ. The utility of an employed worker (index e), who freely chooses working time, takes the form

$$U_e = \left[\theta_C \left(\frac{C_e}{\bar{C}_e} \right)^{\frac{\sigma-1}{\sigma}} + (1 - \theta_C) \left(\frac{F_e}{\bar{F}_e} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \tag{2}$$

where θ_C is the initial value share of consumption, and $\bar{C}$ and $\bar{F}$ are the initial values of consumption and leisure.[2]

U_e is maximised subject to the budget constraint

$$wH_e(1 - t_L^a) = p_C C_e,$$

and the time endowment (T) constraint,

$$F_e + H_e = T.$$

Tax progressivity is captured by a marginal wage tax rate, t_L^m, which in general will differ from the average rate, t_L^a, and has a direct impact on the hours-of-work decision.[3]

---

[2]This is the "calibrated share form" of the CES function (Rutherford, 1998), which simplifies calibration by linking the parameters directly to observable values. In the following, a variable with a bar generally means the value in the initial situation, which is a calibration constant in the counterfactual simulations.

[3]Throughout the paper, the tax schedule is only characterised locally by the average and marginal tax rate. The global form of the schedule (linear progressive, exponential etc.) is left unspecified.

The share parameter in (2) is expressed relative to extended income, Y_E,

$$\theta_C = \frac{\bar{p}_C \bar{C}_e}{\bar{Y}_E}.$$

where extended income includes leisure, valued at the marginal hourly after-tax wage, $w(1 - t_L^m)$:

$$Y_E = w \left[H_e(1 - t_L^a) + F_e(1 - t_L^m) \right] \tag{3}$$

The outcome of the maximisation are demand functions for consumption and leisure:

$$\frac{C_e}{\bar{C}_e} = U_e \left(p_U \frac{\bar{p}_C}{p_C} \right)^\sigma$$

$$\frac{F_e}{\bar{F}_e} = U_e \left(p_U \frac{\bar{w}(1 - \bar{t}_L^m)}{w(1 - t_L^m)} \right)^\sigma,$$

where p_U is the expenditure function corresponding to (2),

$$p_U = \left[\theta_C \left(\frac{p_C}{\bar{p}_C} \right)^{1-\sigma} + (1 - \theta_C) \left(\frac{w(1 - t_L^m)}{\bar{w}(1 - \bar{t}_L^m)} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}$$

Appendices A.1.1 and A.1.2 describe how this function is calibrated to empirical values of labour supply elasticities with respect to income and wage.

2.3 Unemployment

Unemployed workers receive a fixed share of the after tax wage income of the employed as unemployment benefit. This share is given by the replacement rate, c. The unemployed have the same utility function as the employed. However, their consumption and leisure quantities are not determined by optimising choice, but by the demand restriction they face on the labour market. Utility of the unemployed is

$$U_u = \left[\theta_C \left(\frac{C_u}{\bar{C}_e} \right)^{\frac{\sigma-1}{\sigma}} + (1 - \theta_C) \left(\frac{F_u}{\bar{F}_e} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}},$$

where

$$C_u = \frac{cw(1 - t_L^a)H}{p_C}$$

and

$$F_u = (T - \bar{H}) + \delta \bar{H}. \tag{4}$$

7

In (4), I assume that leisure of the unemployed is less that the full time endowment ($\delta < 1$). Economically, this reflects the fact that the unemployed must spend a part of their time for job-seeking and for attending administrative procedures at the social security agency. In any case, involuntary leisure is not likely to generate the same utility as the amount of leisure freely chosen by those with a job. In the calibration, it turns out that for countries with a high replacement rate (see Section 2.9), the utility of the unemployed is higher than for the employed, if time endowment is taken as leisure without discount. As a default, I choose $\delta = 0.5$ in (4). In Section 4.2, this choice will be subject to a sensitivity analysis.

2.4 Participation

Households are assumed homogeneous with respect to their labour-leisure choice, but they differ with respect to their participation decision. This is modelled by heterogeneity in their fixed cost of taking up work, which generates the separation between participating and non-participating individuals. Those with low fixed costs enter the labour market, whereas those with high fixed costs stay at home.[4]

The two step labour-supply decision (participation, hours of work) is solved backwards: First, individuals determine the optimal choice of hours *assuming* that they participate. Second, they compare their fixed cost of working with the outcome of the optimal hours choice, taking the presence of involuntary unemployment into account. In particular, the unemployment-weighted (u) expected utility of supplying labour, U_l, is relevant for the comparison,

$$U_l = (1 - u)U_e + uU_u, \tag{5}$$

which is the same for all individuals. They compare it with their individual fixed cost of supplying labour, U_0, and supply labour if $U_l > U_0$. The calibration of the distribution of U_0 to an empirical participation elasticity is explained in Appendix A.1.3.

---

[4]See Kleven and Kreiner (2006a) for a general discussion of this approach.

2.5 Production

The production sector consists of a large and fixed number, n, of symmetrical firms. Firms are small in the sense that repercussions from their production decisions on the economy-wide aggregate output and price index may be neglected. All firms interact in Dixit-Stiglitz type monopolistic competition (Dixit and Stiglitz, 1977). Each firm faces a demand curve with elasticity η (see Section 2.1).

For firm i, the demand function takes the form

$$x_i = \frac{X}{n}\left(\frac{P}{p_i}\right)^{\eta},$$ (6)

where X and P are communicating aggregate quantity and price indices with

$$X = n^{\frac{1}{1-\eta}}\left(\sum_{i=1}^{n}(x_i)^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}}$$

$$P = \left(\frac{1}{n}\sum_{i=1}^{n}p_i^{1-\eta}\right)^{\frac{1}{1-\eta}},$$

and x_i and p_i are individual quantity and price, respectively. Firm i maximises its profits

$$\pi_i = p_i x_i - w_i L_i$$

(π_i are profits, w_i is the wage and L_i is employment) in p_i, x_i and L_i subject to (6) and the production function

$$x_i = AL_i^{\alpha},$$

treating X, w_i and P as exogenous. Profit maximisation results in the first-order condition

$$p_i = m\frac{w_i L_i}{\alpha x_i},$$ (7)

where m is the mark-up factor ($m = \frac{\eta}{\eta-1}$) applied to marginal output costs (the fraction term on the right-hand side). (7) fixes the income shares at

$$s_L = \frac{w_i L_i}{p_i x_i} = \alpha\left(1 - \frac{1}{\eta}\right)$$

for labour and $s_{\pi} = 1 - s_L$ for profits.

We now consider the symmetrical equilibrium. It is characterised by $p_i = P$, $w_i = w$, $x_i = \frac{X}{n}$ and $L_i = \frac{L}{n}$, for all i. Using (7), this means

$$P = m\frac{wL}{\alpha X}.$$

At this stage, the aggregate producer price level can be normalised: $P \equiv 1$.

2.6   Wage bargaining

Wage formation is modelled as collective bargaining between a trade union and a representative firm. More specifically, I assume (i) that bargaining is only about the wage, not about employment ("right-to-manage" approach)[5], (ii) that the trade union is only concerned with the utility of its employed members ("insider model")[6] and (iii) that hours of work are chosen individually according to the optimisation in Section 2.2, and are not subject of the collective bargain.[7] Wage formation is conceptualised as the maximisation of a Nash function, Ω, where trade unions are represented by the utility mark-up over the fallback option, $U_e - U_a$, and firms by profits, π. The relative bargaining power of the trade union, λ, is an unobservable parameter to be determined in the calibration.

$$\max_{w} \quad \Omega = \left[(U_e - U_a)\right]^{\lambda} \pi \tag{8}$$

The fallback option of the union, U_a, is composed of possible employment in another sector, $\tilde{U}_e$ (with a probability that equals the employment rate), and unemployment (receiving unemployment benefits, see Section 2.3):

$$U_a = (1 - u)\tilde{U}_e + uU_u$$

---

[5]Sørensen (1999) shows that for the type of numerical analysis intended, the choice between right-to-manage and efficient bargaining (where bargaining extends also to the number of employed workers) hardly matters.

[6]Appendix A.1.4 shows that the results are identical to those obtained with a utilitarian union as long as the value shares and the elasticities of labour demand and hours supply are constant.

[7]In the Sørensen (1999) model, it hardly matters quantitatively whether collective bargaining includes hours of work or not. In the model of this paper – with a CES utility function instead of additively separable preferences –, including hours of work in the bargaining set-up would mean a considerable complication of the first-order conditions.

The fallback option is exogenous in the individual wage bargain, so that the first-order condition of the maximisation of the Nash function is given as

$$\lambda \frac{dU_e/dw}{U_e - U_a} + \frac{d\pi/dw}{\pi} = 0$$

Both firms and employed workers make optimal choices, given the wage. This allows us to apply the envelope theorem and express the first-order condition in terms of partial effects:

$$\lambda \frac{\partial U_e/\partial w}{U_e - U_a} + \frac{\partial \pi/\partial w}{\pi} = 0$$

The individual parts of this expression are evaluated as follows, derived from (2):

$$\frac{\partial U_e}{\partial w} = \frac{\partial U_e}{\partial C_e} \frac{\partial C_e}{\partial w} = \theta_C \frac{U_e}{C_e} \frac{H(1 - t_L^m)}{p_C}$$

$$\frac{\partial \pi}{\partial w} = -L$$

where L is total labour input, $L = HN(1-u)$. The first order condition thus becomes

$$\lambda \frac{\theta_C U_e H(1 - t_L^m)}{p_C C_e (U_e - U_a)} - \frac{L}{\pi} = 0 \tag{9}$$

λ is a parameter that cannot directly be observed. Its value is determined in the calibration by inverting (9) and solving for λ, given the values of all other variables in the initial situation (see Appendix A.1.4).

2.7   Public budget

The government collects the following taxes: taxes on wages, profit income, and consumption, with rates t_L^a, t_π and t_C, respectively. The tax revenue is spent for unemployment benefits and public goods. The government budget constraint reads[8]

$$cw(1 - t_L^a)L\frac{u}{1 - u} + p_C C_g = t_L^a wL + t_\pi \Pi + t_C p_C C. \tag{10}$$

In the counterfactual simulations, the exogenous variation is in the marginal wage tax, t_L^m. t_c is kept constant, t_π and t_L^a are endogenously adjusted. The role of t_π is keeping the consumption level of the capitalis household constant, and t_L^a is set so that the amount of public goods, C_g, is held at its initial level. This is "revenue neutral" if we keep in mind that tax revenue is corrected by the expenditure for unemployment benefits.

---

[8]For simplicity, I assume that the consumption tax also applies to government consumption. This does not affect the results.

2.8 Welfare

The welfare criterion used to determine the optimal degree of tax progressivity is the ex-ante expected utility of workers who do not yet know whether they will be employed or unemployed. This is exactly the same indicator that also governs labour supply at the extensive margin (equation 5),

$$U_l = (1 - u)U_e + uU_u.$$

The focus on the utility of workers is justified by the fact that all other magnitudes relevant for a welfare assessment – consumption of the capitalists and the consumptive part of the public budget – are kept at their initial level during the simulations. In the sensitivity analysis, I look also at the case that t_π is fixed and the welfare of capitalists is disregarded (Section 4.3).

2.9 Calibration to OECD economies

The basic model of Section 2 is calibrated to a set of macroeconomic and institutional parameters for a number of OECD countries in 2004/5. The data set contains the six largest European economies: France (FRA), Germany (GER), Great Britain (GBR), Italy (ITA), Spain (ESP) and the Netherlands (NLD), plus the USA and Japan (JPN).

These countries are characterised by seven parameters: share of labour in value added (s_L), average tax rates on consumption (t_C), labour (t_L^a) and capital (t_π), coefficient of residual income progression (CRIP)[9], unemployment rate (u) and replacement rate (c). These parameters are summarised in Table 1 (the exact sources are given in Appendix A.2). In addition, Table 1 reports the unweighted average of the parameters over all countries (row "AVR"), which will be used as a starting point and standard of comparison.

Apart from the country-specific paramters of Table 1, labour supply elasticities are an important input to calibration. These are assumed uniform across countries:

---

[9]The CRIP (coefficient of residual income progression) is defined as the elasticity of after-tax income with respect to pre-tax income, i.e. $(1-t_L^m)/(1-t_L^a)$. In a proportional tax regime, the CRIP is one, and the higher the progressivity of the tax schedule, the lower the CRIP. Jacobsson (1976) derives theoretical properties of this indicator and justifies its use as a measure of tax progressivity.

Table 1: OECD parameters

	s_L	t_C	t_L^a	t_π	CRIP	u	c
FRA	0.601	0.182	0.430	0.266	0.910	0.088	0.630
GER	0.593	0.155	0.392	0.131	0.731	0.079	0.570
GBR	0.654	0.162	0.259	0.352	0.862	0.048	0.340
ITA	0.473	0.151	0.427	0.249	0.844	0.096	0.050
JPN	0.594	0.069	0.269	0.198	0.937	0.050	0.070
NLD	0.582	0.210	0.373	0.212	0.849	0.027	0.600
ESP	0.545	0.142	0.340	0.169	0.878	0.129	0.290
USA	0.627	0.041	0.263	0.246	0.933	0.047	0.060
AVR	0.583	0.139	0.344	0.228	0.868	0.071	0.326

elasticity of hours with respect to the after-tax wage ($\eta_{Hw} = 0.1$), elasticity of hours with respect to income ($\eta_{LY} = -0.1$) and elasticity of participation with respect to the wage ($\eta_{Nw} = 0.2$). The choice of the specific values is motivated in Appendices A.1.1 to A.1.3).

3 Optimal tax progressivity

In this section, I perform numerical simulations to determine the optimal degree of tax progressivity and identify its driving forces. In Section 3.1, I explain the determination of the optimal tax progressivity in the "average OECD" model. Then the driving forces are identified in two sets of numerical exercises. First, starting from the average values, I vary one parameter at a time to obtain partial effects on optimal tax progressivity. Second, I run the fully specified country models and check to what extent the deviations from the average-OECD outcome can be decomposed into effects of variations in the individual parameters.

3.1 "Average OECD" model

As a point of reference, I calibrate the basic model of Section 2 to the unweighted OECD averages for all country-specific parameters (row "AVR" of Table 1). In this model version, I perform a number of tax variations that allow us to develop a feeling for the range of tax rates that we expect to be relevant in the model runs to follow. They are summarised in Figure 1, which shows characteristic tax rate constellations in the (t_L^m, t_L^a) space.

Figure 1: Average OECD model: characteristic tax structures
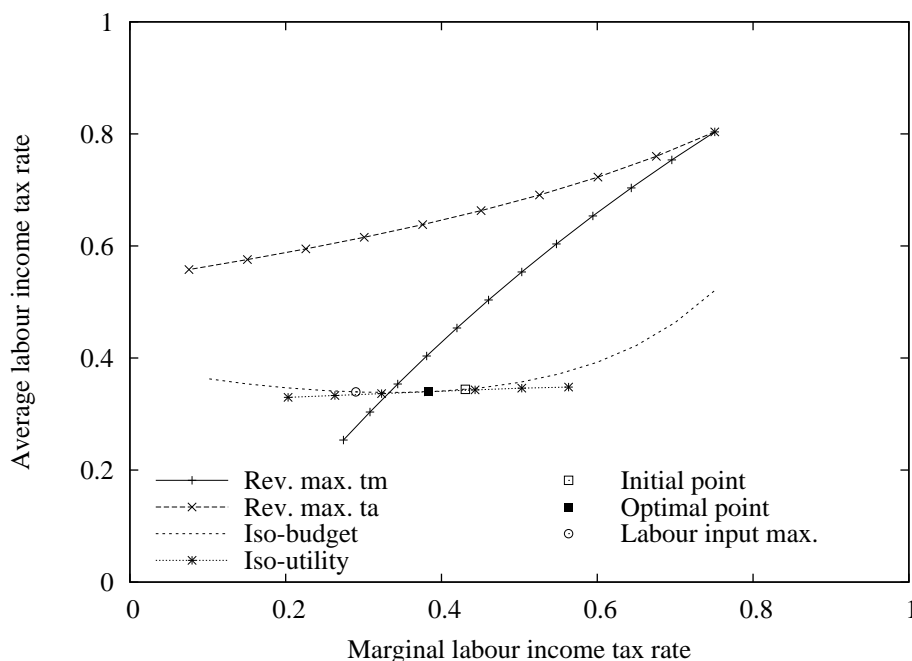
The upper rightmost point of Figure 1, where the curves "Rev. max. tm" and "Rev. max. ta" meet, is the point of maximum tax revenue.[10] We may call it "Leviathan point", because it is the point a malevolent, exploitative dictator would choose ($t_L^a = 80\%$, $t_L^m = 75\%$). The curves "Rev. max. tm" and "Rev. max. ta" connect the points of *partial* revenue maxima. "Rev. max. tm" gives the points of maximum tax revenue when t_L^m is varied and t_L^a is held fixed at its respective value,

---

[10]To be precise: It is the point where tax revenue is maximised through simultaneous variation of the labour tax parameters t_L^a and t_L^m, taking all general equilibrium interactions into account, compensating the capitalist by an adjustment of t_π, and treating the value of t_C as given.

14

i.e. the revenue maximum on each horizontal line through the tax space. "Rev. max. ta" is the same for t_L^a and t_L^m exchanged, i.e. the maximum on a vertical line.[11] The dot "Initial point" marks our point of departure, $t_L^a = 34.4\%$, $t_L^m = 43.1\%$ (CRIP= 0.87). Through this point, we draw an iso-budget line ("Iso-budget"). To the left of the curve of revenue-maximising t_L^m rates, the iso-budget line will be downward-sloping. In terms of tax revenue, higher levels of t_L^a are traded off against higher levels of t_L^m. At the intersection point, the iso-budget line is horizontal, because a further increase of t_L^m does not lead to additional tax revenue any more, which would allow the tax planner to lower t_L^a as a compensation.[12] Consequently, the initial tax revenue cannot be materialised with t_L^a below 33.9%.

The iso-budget line summarises the set of choice options for the optimal taxation problem (at given levels of public goods and consumption of the capitalists). It remains to be determined which of these options should be chosen. This is captured by the "Iso-utility" line, which connects points of the same expected utility of workers, taking into account all general equilibrium interactions (i.e. adjusting wages and unemployment rates). It turns out that the iso-utility line is slightly increasing at its tangency point with the iso-budget line, which is in the upward-sloping range between the minimum for t_L^a and the initial point. The optimal t_L^m is 38.3%, almost five percentage points lower than the initial 43.1%. This gives an optimal CRIP of 0.93, compared to the initial level of 0.87.[13] Finally, Figure 1 also shows the point where total labour input (product of individuals and hours) is maximised on the iso-budget line. This is considerably far from the utility maximum at a marginal tax rate of 29.0% and a CRIP of 1.08, i.e. a *regressive* tax. The reason for this difference is that additional labour input at the intensive margin is relatively cheap in welfare terms, while unemployment is expensive. Therefore reducing unemployment (by higher tax progressivity) is welfare-enhancing even if this does not compensate the labour volume loss caused by a decrease in hours of work.

---

[11]Interestingly, these curves have a positive slope. Usually, two taxes on different goods or factors of production result in negatively sloped curves. See the figures in Boeters (2004).

[12]Analogously, the iso-budget curve would be vertical when it meets the "Rev. max. ta" line. However, in the constellation of Figure 1 this point is not reached with positive levels of t_L^m.

[13]This considerably deviates from the optimal CRIP level of 0.72 that Sørensen (1999) obtains in his model. Additional model runs showed that the income elasticity of labour supply (which is zero in the Sørensen (1999) model) is the most likely candidate for an explanation of this discrepancy.

An alternative way of visualising the optimal tax problem is by depicting the marginal cost of public funds (MCPF) of the taxes involved.[14] This is shown in Figure 2. Here we see that the MCPF of t_L^a is quite stable at a level of about 2, which indicates a considerable excess burden of taxation. The MCPF of t_L^m, on the other hand, behaves rather atypically. On the downward-sloping part of the iso-budget line, it is negative, because increasing t_L^m causes both tax revenue and private welfare to rise. At the lowest point of the iso-budget line at about $t_L^m = 33\%$, it has a singularity, because the marginal tax revenue is zero. To the right of this point, MCPF of t_L^m is positive, which results from the combination of negative marginal tax revenue and positive effects on private welfare. This is also the region where both MCPF curves intersect and we have the welfare maximum at $t_L^m = 38.3\%$. Finally, for very high levels of t_L^m, MCPF again turns slightly negative, because now the effects on private welfare are also negative.

Figure 2: Average OECD model: Marginal cost of public funds

---
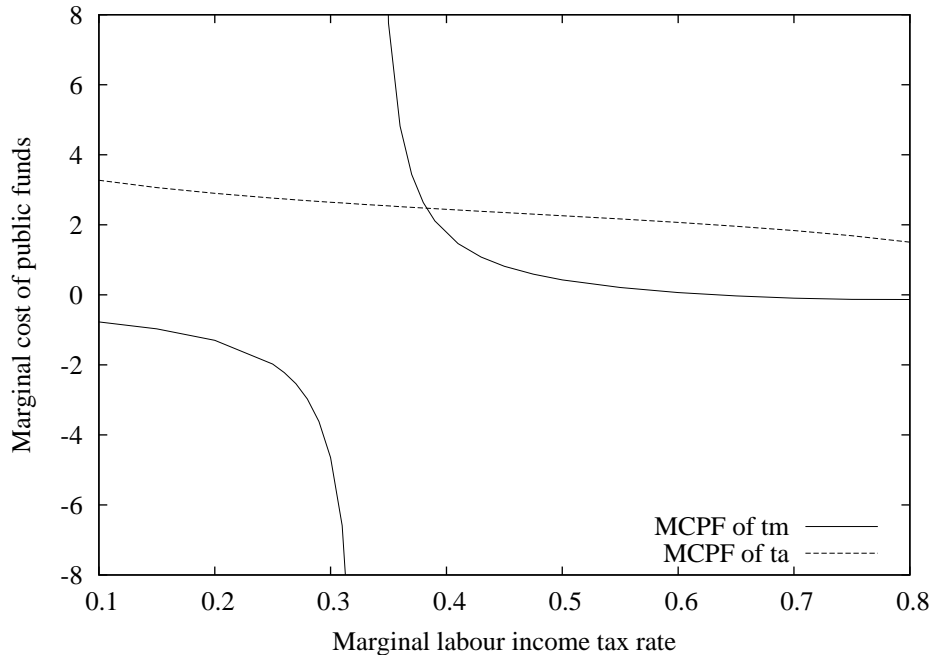
[14]The MCPF has been numerically calculated by using different starting points on the iso-budget line of Figure 1, changing the respective tax by a small amount, and dividing the negative monetarised change in expected utility by the additional tax revenue.

3.2 Systematic parameter variation

Starting from the average OECD parametrisation of Section 3.1, the parameters that lead to country heterogeneity are now varied systematically. In order to isolate the effects of the different parameters, I replace, one by one, the average parameter value with both its minimum and maximum value in all countries considered. This yields 14 model variations, two for each of the seven parameters. The results of these model exercises are summarised in Table 2.

Table 2: Systematic parameter variation

| | minimum | | maximum | | |
| | | optimal | | optimal | partial |
	value	CRIP	value	CRIP	effect
s_L	0.473 (ITA)	0.934	0.654 (GBR)	0.935	0.002
t_C	0.041 (USA)	0.915	0.210 (NLD)	0.947	0.188
t_L^a	0.259 (GBR)	0.907	0.430 (FRA)	0.961	0.316
t_π	0.131 (GER)	0.934	0.352 (GBR)	0.934	0.000
CRIP	0.731 (GER)	0.802	0.937 (JPN)	0.998	0.952
u	0.027 (NLD)	1.180	0.129 (ESP)	0.777	-3.944
c	0.050 (ITA)	0.946	0.630 (FRA)	0.839	-0.184

Table 2 shows the minimum and maximum value of each of the parameters in the dataset, the respective country and the resulting optimal degree of tax progressivity. For the interpolation of the country results, I assume that the effects of the individual parameters are linear and additive. This allows us to calculate partial effects, which are shown in the last column of Table 2. Take the partial effect for t_C (0.188) as an example. It may be interpreted in the following way: A one-percentage-point increase in the consumption tax rate leads to a value of the optimal CRIP that is about 0.2 percentage point higher. The spread between the optimal CRIP for the minimum and maximum value of a coefficient is informative with respect to the relevance of the respective coefficient for explaining country differences. A parameter is particularly relevant if there actually is variation between countries *and* its partial effect is high.

Interpreted in this way, Table 2 shows that most "action" is in the CRIP and the unemployment rate.[15]

3.3 Explaining cross-country differences

In which way can the results from the systematic parameter variation be used to explain cross-country differences? Table 3 shows the actual and the optimal CRIP for all countries. Row "deviation" reports the difference between the optimal country-specific CRIP and the optimum in the average OECD model (0.93). This difference is "explained" by the deviations of all parameters from their OECD averages, multiplied with the partial effects from Table 2.

Table 3: Cross-country differences

	FRA	GER	GBR	ITA	JPN	NLD	ESP	USA
CRIP								
actual	0.910	0.731	0.862	0.844	0.937	0.849	0.878	0.933
optimal	0.838	0.747	0.999	0.869	1.041	1.093	0.789	1.038
deviation[a]	-0.096	-0.188	0.065	-0.066	0.107	0.158	-0.145	0.103
Partial effects								
s_L	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
t_C	0.008	0.003	0.004	0.002	-0.013	0.013	0.001	-0.018
t_L^a	0.027	0.015	-0.027	0.026	-0.024	0.009	-0.001	-0.026
t_π	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
CRIP	0.040	-0.130	-0.006	-0.023	0.065	-0.018	0.009	0.062
u	-0.069	-0.034	0.089	-0.101	0.081	0.172	-0.231	0.093
c	-0.056	-0.045	-0.003	0.051	0.047	-0.050	0.007	0.049
Sum	-0.050	-0.190	0.058	-0.044	0.157	0.125	-0.215	0.160

[a]Deviation from the optimal CRIP in the "OECD average" model (0.934).

The fit is not perfect, but fairly good. The R^2 of this exercise (when interpreting, without any statistical implication, the partial effects as regression coefficients) is

---

[15]The surprising fact that actual initial progressivity has a strong effect on optimal progressivity is discussed in Sections 3.4 and 4.1.

0.88, which means that the assumption of additive, linear effects is justified. As it stands, the regression may be read in the following way: Low optimal progressivity in the Netherlands and Great Britain is mostly driven by low unemployment rates. Reversely, high unemployment in Spain makes a high degree of tax progressivity desirable. In contrast to these countries, high optimal tax progressivity in Germany is mostly driven by the low initial CRIP. Low optimal tax progressivity in the USA and Japan is explained by a combination of relatively low unemployment and a high initial CRIP.

The decomposition results thus confirm what we saw in Table 2. The largest contribution to the explanation of the effects comes from the unemployment rate and the CRIP. The contributions of the replacement rate, the consumption tax and the average wage tax are smaller. The value share of labour and the profit tax are negligible.

3.4 Interpretation of the regression coefficients

Until this point, the analysis has been mainly descriptive. It can be shown by systematic parameter variation that differences in the degree of optimal tax progressivity between countries are to a large extent driven by variations in the unemployment rate and initial tax progressivity. But also the level of other taxes and the replacement rate play a role. How are these effects explained economically?

Most straightforward is the effect of unemployment on optimal tax progressivity. Recall that the positive welfare effects of higher tax progressivity are driven by the wage-depressing and unemployment-reducing forces of higher tax progressivity in wage bargaining. The higher the initial unemployment level, the higher (in percentage points) is the reduction in unemployment through a given increase in tax progressivity, and the more people benefit from this by switching into employment. Therefore, countries with a high unemployment level have higher optimal tax progressivity. In the opposite extreme case, if there is no (or very low) unemployment, there is nothing to be gained from increasing tax progressivity, while there is still the distortionary effect on labour supply.

If we want to explain the effect of the existing levels of t_L^a, and t_C on optimal tax progressivity[16] (the higher these levels, the lower optimal tax progressivity), we need an intermediate step. It is important to know that choosing the utility-maximising point on the iso-budget line in Figure 1 is quite distinct from maximising the total labour input.[17] The labour input maximum is considerably farther to the left of the utility-maximising point (at $t_L^m = 29.0\%$). In moving from here to the right, labour input decreases, i.e., the employment of additional workers (reduction in unemployment) is over-compensated by the loss in hours per person employed. This is in fact welfare-enhancing, because the employed are much closer to their labour-leisure optimum than the unemployed. Thus, an hour of work of someone formerly unemployed is worth more than an hour lost of someone who was and remains employed. However, this trade-off is altered by the existing taxes. If there are taxes, a loss in employment also means a loss in tax revenue, which must be compensated if the government's budget is to remain balanced. Therefore, the volume of labour input carries more weight in the welfare trade-off if taxes are high. The higher the taxes, the closer we remain to the labour input maximum, i.e., the less we move towards higher tax progressivity.

The effect of the replacement rate takes place via the public budget as well. A high replacement rate means high budgetary costs of unemployment. Reducing unemployment thus becomes more attractive. We have an additional positive effect on the public budget, which translates into compensatory adjustments of t_L^a, which modifies the trade-off in the direction of higher tax progressivity.

We are left with the effect of the initial CRIP on the optimal CRIP, which is large and positive. An increase in the initial CRIP translates almost one-to-one into an increase in the optimal CRIP. This is strange at first sight. How is it possible that the initial level of one variable determines its own optimal value? If we want to understand this effect, we need to go back to the calibration of the utility function to the elasticity of labour supply at the intensive margin (see Appendix A.1.2). For given (distribution and substitution) parameters of the utility function, an increase in tax progressivity that the agent faces would lead to a higher labour supply elasti-

---

[16]The level of t_π has no effect, because it is adjusted in the simulations.

[17]In the model of this paper, the labour input maximum coincides with the output maximum and the maximum of tax revenue from the consumption tax.

city. For this very reason the marginal wage tax rate becomes the more distortive the higher it is. However, we calibrate utility functions in all countries to the same wage elasticity of labour supply. This means that for countries with high initial progressivity (used for calibration), this effect must be counterbalanced by parameters that work in the direction of a low elasticity (namely, a low elasticity of substitution between consumption and leisure). Controlled for tax progressivity, labour supply is therefore less elastic in countries with high initial tax progressivity. This in turn means that the negative effects of higher tax progressivity are less severe, so that optimal tax progressivity follows actual progressivity.

The argument in the previous paragraph allows us to trace back the effect to the calibration of labour supply and makes it seem less strange than it was in the beginning. Most importantly, we must not interpret the effect as *causal*. Of course, governments cannot change the level of the optimal tax progressivity just by changing progressivity itself. It is plausible, however, to interpret the effect as *indicative*. If it is the case that in certain countries the same level of labour supply elasticity is reached only at higher degrees of tax progressivity, governments can exploit this fact in setting the tax schedule. Politicians find by trial and error that making taxes more progressive is not too distorting. This leads to higher tax progressivity in these countries, which in turn drives labour supply elasticities up to a level which is close to the average of other countries. The regression results suggest that governments are relatively good in implicitly adjusting the tax rates to the specific labour market conditions they face.

However, against the background of the results in this section, we must reconsider the question of whether it is a good idea to assume identical labour supply elasticities in all countries. Section 4.1 explores an alternative way of calibrating the utility functions, which allows for differences in the labour supply elasticities across countries.

4   Sensitivity analysis

In this section, I check the model results for robustness by exploring three themes that have emerged in the discussion so far. Is it possible to decouple the actual and

the optimal level of tax progressivity by changing the calibration of labour supply? (Section 4.1) To what extent do the results depend on the choice of the (arbitrary) parameter that determines the valuation of involuntary leisure? (Section 4.2) What effect has the adjustment of the profit tax rate on the results? (Section 4.3)

4.1 Calibration of labour supply elasticity

The fact that initial tax progressivity plays an important role in determining optimal progressivity, is explainable, but needs a closer inspection. This was the conclusion at the end of Section 3.4. The critical step that leads to this result is calibrating the utility function in all countries to the same values of the labour supply elasticities at the intensive margin (with respect to income, η_{LY}, and with respect to the wage, η_{Hw}). These exogenous elasticity values can only be combined with different levels of the tax rates, t_L^a and t_L^m, if the parameters of the utility function, θ_C, σ and T, are allowed to vary across countries. This is shown in the left hand panel of Table 4. The calibrated elasticity of substitution between consumption and leisure varies between 1.73 and 1.93. It moves in parallel with the CRIP (see Table 1).[18] So countries with low initial tax progressivity have a low elasticity of substitution, which weakens labour supply responses at the intensive margin and makes tax progressivity more attractive. This is exactly what produces the large regression coefficient for the CRIP in Section 3.2.

What if we instead treated σ as the deeper parameter that does not change across countries? The right hand panel of Table 4 shows this case. To arrive at the elasticity values displayed, I have taken the average OECD model as point of departure, and then – with fixed σ and θ_C – implemented the country-specific parameters[19] as a counterfactual. The resulting elasticity values are displayed in columns "η_{LY}" and "η_{Hw}". We see that high tax progressivity (GER) now leads to high absolute values of the elasticities, and conversely (low progressivity in JPN and USA). The values of η_{LY} and η_{Hw} from Table 4 are in turn used to re-calibrate the model at the initial point. This leads to a slight adjustment of θ_C, because the value share of leisure is

---

[18]With the elasticity values chosen as exogenous, σ happens to be precisely $1 + \text{CRIP}$.

[19]The only exception is the unemployment rate, which cannot be imposed on the model as an exogenous parameter.

not constant with a CES function.[20]

Table 4: Calibration of labour supply

| | η_{LY}, η_{Hw} exogenous | | | σ exogenous | | | |
	θ_C	σ	T	θ_C	η_{LY}	η_{Hw}	T
FRA	0.909	1.910	1.110	0.911	-0.098	0.094	1.107
GER	0.927	1.731	1.108	0.898	-0.140	0.159	1.156
GBR	0.914	1.862	1.109	0.925	-0.086	0.087	1.093
ITA	0.916	1.844	1.109	0.886	-0.135	0.138	1.152
JPN	0.906	1.937	1.110	0.928	-0.077	0.071	1.083
NLD	0.915	1.849	1.109	0.908	-0.109	0.111	1.120
ESP	0.912	1.878	1.110	0.910	-0.103	0.101	1.113
USA	0.907	1.933	1.110	0.933	-0.072	0.067	1.077
AVR	0.913	1.868	1.110	0.913	-0.100	0.100	1.110

If we repeat the exercise of Sections 3.2 and 3.3 with the alternative method of labour supply calibration, we obtain the regression coefficients listed in Table 5.

Comparing Tables 5 and 2, we see that the coefficient of the CRIP is actually reduced to less than half of its former value. It remains positive, however. In compensation, the coefficients of t_C and t_L^a have increased substantially and the coefficient of s_L has turned negative. This is reflected in the revised decomposition in Table 6. For the CRIP, the explanatory contribution is significantly reduced, while it remains high for the unemployment rate and increases for t_L^a.

The fact that the coefficient of the CRIP is still positive comes unexpected, given that the set-up of this section was aimed at correcting for this effect. The most probable explanation is based on the interaction between tax progressivity and unemployment. Higher initial tax progressivity induces lower initial unemployment.

---

[20]Unfortunately, the clear-cut distinction between calibration and simulation breaks down at this point. The different values of η_{LY} and η_{Hw} do not only depend on t_L^a and t_L^m, but also on the deviation of H from its initial value. This, in turn, depends on the wage and is only determined in full equilibrium.

Table 5: Partial effects with alternative labour supply calibration

	minimum		maximum		
	value	optimal CRIP	value	optimal CRIP	partial effect
s_L	0.473 (ITA)	0.963	0.654 (GBR)	0.917	-0.254
t_C	0.041 (USA)	0.901	0.210 (NLD)	0.956	0.328
t_L^a	0.259 (GBR)	0.889	0.430 (FRA)	0.985	0.560
t_π	0.131 (GER)	0.934	0.352 (GBR)	0.934	0.000
CRIP	0.731 (GER)	0.877	0.937 (JPN)	0.959	0.399
u	0.027 (NLD)	1.180	0.129 (ESP)	0.777	-3.944
c	0.050 (ITA)	0.948	0.630 (FRA)	0.835	-0.195

As unemployment is a separate regressor, the total effect of initial tax progressivity is decomposed into the (positive) direct effect captured by the parameter in Table 5 and the (negative) indirect effect via unemployment. Added up, these two effects might well be zero. However, performing such a check would require a separation of tax-progressivity-induced unemployment from unemployment caused by other country-specific conditions. I have not succeeded in finding a separation method that is consistent with the calibration of the wage bargaining system (Appendix A.1.4) and leave it therefore at this qualitative discussion.

Without further empirical analysis of differences in labour supply elasticities across countries, it cannot be decided which of the two calibration variants is to be preferred. At a first glance, having initial tax progressivity determining optimal progressivity seems dubious. However, given the results of this section, the argument from the end of Section 3.4 can be reinforced: If it is really the case that countries with higher tax progressivity have the same labour supply elasticities as countries with low progressivity, the effect is not unreasonable. It must be interpreted as an indicative rather than a causal effect, however. Cross-country differences in tax progressivity reflect differences in labour supply conditions that would surface if taxes were the same everywhere.

Table 6: Cross-country differences with alternative labour supply calibration

	FRA	GER	GBR	ITA	JPN	NLD	ESP	USA
CRIP								
actual	0.910	0.731	0.862	0.844	0.937	0.849	0.878	0.933
optimal	0.834	0.800	0.969	0.946	0.976	1.097	0.793	0.961
deviation[a]	-0.100	-0.134	0.035	0.011	0.041	0.162	-0.141	0.027
Partial effects								
s_L	-0.005	-0.002	-0.018	0.028	-0.003	0.000	0.010	-0.011
t_C	0.014	0.005	0.008	0.004	-0.023	0.023	0.001	-0.032
t_L	0.048	0.027	-0.048	0.047	-0.042	0.016	-0.002	-0.046
t_π	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
CRIP	0.017	-0.055	-0.002	-0.009	0.027	-0.008	0.004	0.026
u	-0.069	-0.034	0.089	-0.101	0.081	0.172	-0.231	0.093
c	-0.059	-0.047	-0.003	0.054	0.050	-0.053	0.007	0.052
Sum	0.054	-0.106	0.026	0.022	0.091	0.150	-0.211	0.082

[a]Deviation from the optimal CRIP in the "OECD average" model (0.934).

To my knowledge there is no example of a meta-analysis of labour supply elasticities in the literature that takes tax progressivity as a regressor.[21] We thus remain agnostic about the relative performance of the two calibration methods. All we can say is that the sign of the effects does not depend on the approach chosen.[22]

4.2 Utility from involuntary leisure

In Section 2.3, we have seen that we are left with one free parameter in the calibration, which cannot be empirically founded: the utility from involuntary leisure, δ. The value of δ cannot be one, because then – given the specification of utility as a CES function of consumption and leisure and given empirically plausible values of the elasticities of labour supply – the unemployed would be better off than the

---

[21]Evers et al. (2005), e.g., has no regressors that capture institutions of the countries the studies reviewed are about.

[22]The only exception is the value share of labour, but this effect is small, anyway.

25

employed.[23] In the light of the discussion on a "poverty trap", this might even be realistic at least for a subset of individuals. However, it is not consistent with the wage bargaining set-up, which is based on the assumption that trade unions bargain over a wage that gives the employed utility *on top of* what they would have if they were unemployed.

In Section 2.3, I apply an ad-hoc solution to this problem by assuming that involuntary leisure gives only half the utility of leisure chosen voluntarily ($\delta = 0.5$). Now I investigate to what extent the results are sensitive to this assumption. As an alternative, extreme choice, I set the utility-of-involuntary-leisure parameter to zero ($\delta = 0$). We might think of an unemployment benefit system that requires benefit recipients to work full-time in some public services. The consequences of this change are listed in Table 7 (with the same set-up as Table 2).

Table 7: Partial effects with utility of involuntary leisure set to zero

| | minimum | | maximum | | |
	value	optimal CRIP	value	optimal CRIP	partial effect
s_L	0.473 (ITA)	0.922	0.654 (GBR)	0.922	0.002
t_C	0.041 (USA)	0.900	0.210 (NLD)	0.935	0.204
t_L^a	0.259 (GBR)	0.892	0.430 (FRA)	0.951	0.343
t_π	0.131 (GER)	0.922	0.352 (GBR)	0.922	0.000
CRIP	0.731 (GER)	0.791	0.937 (JPN)	0.984	0.939
u	0.027 (NLD)	1.166	0.129 (ESP)	0.764	-3.940
c	0.050 (ITA)	0.936	0.630 (FRA)	0.879	-0.098

Comparing Tables 2 and 7 shows that optimal tax progressivity is higher (lower CRIP) across the board.[24] This is a plausible outcome. Through the lower valuation of involuntarily leisure utility differences between the employed and the unemployed

---

[23]Strictly speaking, this only applies to countries with a high replacement rate.

[24]There is a single exception from this general picture: The optimal tax progressivity for the maximum value of the replacement rate, c, goes down instead of up. The mechanism is intricate. If the replacement rate is high, the most important mechanism is that unemployment affects the

become larger. Hence larger gains in expected utility are achieved by reducing unemployment. And higher tax progressivity is the means to generate this outcome.

The change in the level of optimal tax progressivity is almost uniform across all model variations. This means that the marginal effects, captured by the regression coefficients, remain almost the same as in the main variant of the model. The unemployment rate and initial tax progressivity remain the most important determinants of optimal tax progressivity.

4.3   No adjustment of the profit tax

In the counterfactual simulations of Section 3, we have assumed that the profit tax is adjusted so as to keep the consumption level of the capitalists at a constant level. This allows us to restrict ourselves on the expected utility of workers in the welfare analysis, because these are the only households that experience welfare changes. However, a coordinated adjustment of profit and labour income taxes is not particular realistic as a policy scenario (at least not in the way assumed in the simulations: that capitalists are compensated for the indirect effects from the labour tax reform). To put the simulations into perspective, I run them once again without the adjustment of the profit tax, and disregarding the income changes of the capitalists. This can be given two different interpretations: either all capital is owned by foreigners, whose welfare does not enter the target function of the domestic government. Or we consider the case of a "labourist" government, which only cares about the welfare of workers.

Table 8 shows the results of the decomposition analysis. Disregarding welfare changes for the capitalist has a general level effect on optimal tax progressivity: Optimal tax progressivity is higher than in the base case, because capitalist need not to be compensated for the income losses they suffer when we move away from the labour-input-maximising point in Figure 1. In the "average OECD" model, the

_____

utility of the employed through the adjustments of t_L^a necessary to compensate the public budget for unemployment benefit payments. It turns out that with lower valuation of unemployed time, the reductions in unemployment through higher tax progressivity are slightly lower than in the main variant. This pushes the effect in the opposite direction.

optimal CRIP is now at 0.85 (compared to 0.93 in the base case with compensation). Interestingly, this is almost exactly equal to the actual CRIP of 0.87.

Table 8: Systematic parameter variation

| | minimum | | maximum | | |
| | | optimal | | optimal | partial |
	value	CRIP	value	CRIP	effect
s_L	0.473 (ITA)	0.819	0.654 (GBR)	0.870	0.286
t_C	0.041 (USA)	0.812	0.210 (NLD)	0.875	0.371
t_L^a	0.259 (GBR)	0.814	0.430 (FRA)	0.890	0.443
t_π	0.131 (GER)	0.837	0.352 (GBR)	0.869	0.147
CRIP	0.731 (GER)	0.742	0.937 (JPN)	0.903	0.780
u	0.027 (NLD)	1.036	0.129 (ESP)	0.722	-3.075
c	0.050 (ITA)	0.855	0.630 (FRA)	0.786	-0.119

The partial effects of Table 8 are in a similar range to the base case (Table 2). Some quantitative changes need comment, however. The coefficient of the value share of labour (s_L) is now positive. This is plausible, because the higher the share of labour, the lower profits and the smaller the external effect on capitalists, which works in favour of high tax progressivity. The partial effects of the tax rates (t_C, t_L^a, and, for the first time, t_π itself) are stronger. This is a tax-base effect. We are farther away from the labour-input maximum. The tax base for the wage tax is therefore smaller. A given tax revenue loss from shrinking economic activity (which is the higher the higher the tax rates) now requires a larger adjustment of the average wage tax. This makes additional tax progressivity, which produces tax revenue losses, less attractive.

Table 9 shows the decomposition for the individual countries. s_L and t_π now enter the picture with contributions to the explanation of deviations from the OECD average. However, as their coefficients are not particularly large and cross-country variation is moderate for these variables, the overall qualitative pattern remains unchanged.

Table 9: Cross-country differences

	FRA	GER	GBR	ITA	JPN	NLD	ESP	USA
CRIP								
actual	0.910	0.731	0.862	0.844	0.937	0.849	0.878	0.933
optimal	0.807	0.700	0.930	0.786	0.883	1.006	0.717	0.895
deviation[a]	-0.045	-0.152	0.078	-0.066	0.031	0.154	-0.135	0.043
Partial effects								
s_L	0.005	0.003	0.020	-0.031	0.003	-0.001	-0.011	0.012
t_C	0.016	0.006	0.009	0.005	-0.026	0.026	0.001	-0.036
t_L^a	0.038	0.021	-0.038	0.037	-0.033	0.013	-0.002	-0.036
t_π	0.006	-0.014	0.018	0.003	-0.004	-0.002	-0.009	0.003
CRIP	0.032	-0.107	-0.005	-0.019	0.054	-0.015	0.008	0.051
u	-0.054	-0.026	0.069	-0.078	0.063	0.134	-0.180	0.072
c	-0.036	-0.029	-0.002	0.033	0.030	-0.033	0.004	0.032
Sum	0.007	-0.146	0.072	-0.051	0.087	0.122	-0.188	0.097

[a]Deviation from the optimal CRIP in the "OECD average" model without t_π compensation (0.852).

5   Conclusions

The model of the paper is used to analyse the determinants of optimal tax progressivity in a labour market with collective wage bargaining and flexible labour supply. The framework chosen is simple enough to keep an overview over the basic mechanisms, but has sufficient complexity to be calibrated to a number of behavioural and macroeconomic parameters, partly universal, partly country-specific.

In particular, the following parameters are used to adjust the model to country-specific conditions: factor shares in value added, important macroeconomic tax quotas (consumption tax, labour tax, capital tax), unemployment rates, tax progressivity and replacement rate. In addition, three behavioural parameters were taken into account: wage elasticity of labour supply at the intensive and at the extensive margin, and income elasticity of labour supply.

When calibrated to average OECD values, the model produces the following results:

- Optimal tax progressivity is lower than actual progressivity. Moving to the optimal point would mean a five percent cut in the marginal wage tax, which would allow for half a percent compensating decrease in the average wage tax.

- The optimal tax structure is located in the region where increasing the marginal labour tax produces negative tax revenue.

- The optimal tax structure is considerably far away from the point of maximum labour input. This is because an labour input gain through lower unemployment has more value (in welfare terms) than labour input losses at the hours-of-work margin.

There is considerable cross-country variation both in actual and optimal tax progressivity. A decomposition approach shows that approximating the general effect by linear additive effects of the seven country-specific parameters gives a reasonable fit, where:

- Higher initial unemployment leads to higher optimal tax progressivity, because unemployment reduction effects are approximately proportional to the initial level.

- A higher general tax level leads to lower optimal tax progressivity, because the tax revenue losses when departing from the employment maximum, which must be compensated through a higher average labour tax, are higher.

- Initial tax progressivity affects optimal progressivity. The extent of this effect depends on the exogeneity assumptions in the calibration. If labour supply elasticities are assumed to be exogenous and constant across countries, the effect it large. If the elasticity of substitution between consumption and leisure is considered a "deep", constant parameter, the effect is considerably smaller.

This last point may be translated into a first indication of where the results may be refined in future work. A careful meta-study on labour supply elasticities could

explore which parameters of an underlying utility function should most reasonably be considered to be "deep" constants. Such an approach could re-establish the calibration-simulation dichotomy, which becomes blurred in the alternative labour supply calibration of Section 4.1.

There are other points where the analysis of this paper potentially can be improved upon. The wage bargaining model used does not have a sufficient number of parameters[25] to be calibrated to empirical estimates of wage equations. There are problems at the other end of the model-empirics match as well. Empirical wage curve estimates turn out to be very unstable, so that it is difficult to find a good standard of comparison (Folmer, 2009). Nevertheless, it could be illuminating to replace the wage bargaining equation in the model with wage curves that are more similar to empirical specifications. Other aspects that could be integrated in the model in a straightforward manner are interactions between different skill types of workers and international capital mobility, which both are likely to have an impact on the results. However, here we approach the grey area of "real" applied models. These focus on institutional detail at the cost of more and more intertwined economic effects that can only be disentangled with great effort. For a recent example of such a model with different skill types, labour supply reactions at the micro level, a differentiated sectoral structure and mobile capital, see Boeters and Feil (2009).

Finally, if we put aside all doubts about the working mechanisms of the model and take the levels of optimal tax progressivity that we obtain at face value, we may proceed to a question of political economy: What might *explain* the deviations of actual tax progressivity from the level characterised as "optimal" in this study? Given that there are deviations in both directions, this is not an easy question to answer. The results of the paper apply to a representative ex-ante worker and do not capture distributional considerations. Governments with large redistributive ambitions are likely to choose more progressive tax schedules than the "optimal" ones of this paper. On the other hand, if the actual tax policy in a country is determined by labour market insiders, whose unemployment risk is lower than the average unemployment rate, this would lead to less tax progressivity than the reference values

---

[25]Actually, there is one free parameter, the value of involuntarily unemployed time. The effect of this parameter on the results is small (see Section 4.2), so that it cannot be used to tune the bargaining equations to empirically estimated wage curve elasticities.

of this paper. It might be an interesting line of research to see whether empirical indicators of redistribution willingness or insider power in unions are better correlated with deviations between actual and optimal tax progressivity than with actual tax progressivity itself. But considering the small size of the potential dataset and the many additional assumptions that enter the calculations of the optimal tax rates, it would be a surprise if such an empirical analysis produced strong, significant results.

References

Aaberge, Rolf, Ugo Colombino, Erling Holmøy, Birger Strøm and Tom Wennemo (2004), Population ageing and fiscal sustainability: an integrated micro-macro analysis of required tax changes, Statistics Norway Discussion Paper 367.

Arntz, Melanie, Stefan Boeters, Nicole Gürtzgen and Stefanie Schubert (2008), Analysing welfare reform in a microsimulation-AGE model: The value of disaggregation, Economic Modelling 25, 422–439.

Ballard, Charles L. (2000), How many hours are in a simulated day? the effects of time endowment on the results of tax-policy simulation models, Discussion Paper, Michigan State University, URL `http://www.econ.msu.edu/faculty/ballard/HoursPaper.pdf`.

Boeters, Stefan (2004), Green tax reform and employment: The interaction of profit and factor taxes, FinanzArchiv 60, 222–239.

Boeters, Stefan and Michael Feil (2009), Heterogeneous labour markets in a microsimulation-AGE model: Application to welfare reform in Germany, Computational Economics 33.

Bovenberg, A. Lans, Johan J. Graafland and Ruud A. de Mooij (2000), Tax reform and the Dutch labor market: An applied general equilibrium approach, Journal of Public Economics 78, 193–214.

Dixit, Avinash and Joseph E. Stiglitz (1977), Monopolistic competition and optimum product diversity, American Economic Review 67, 297–308.

Evers, Michiel, Ruud A. de Mooij and Daniel van Vuuren (2005), What explains the variation in estimates of labour supply elasticities?, CESifo Working Paper 1633.

Folmer, Kees (2009), Why do macro wage elasticities diverge? A meta analysis, CPB Discussion Paper 122.

Fuest, Clemens and Bernd Huber (1998), Tax progression and human capital in imperfect labour markets, EPRU Discussion Paper 1998-03.

Graafland, Johan J., Ruud A. de Mooij, André G.H. Nibbelink and Ate Nieuwenhuis (2001), MIMICing tax policies and the labour market, Elsevier, Amsterdam.

Hersoug, Tor (1984), Union wage responses to tax changes, Oxford Economic Papers 36, 37–51.

Holmlund, Bertil and Ann-Sofie Kolm (1995), Progressive taxation, wage setting and unemployment: Theory and Swedish evidence, Swedish Economic Policy Review 2, 423–460.

Jacobsson, Ulf (1976), On the measurement of the degree of tax progression, Journal of Public Economics 5, 161–168.

Kleven, Henrik Jacobsen and Claus Thustrup Kreiner (2006a), The marginal cost of public funds: Hours of work versus labor force participation, Journal of Public Economics 90, 1955–1973.

Kleven, Henrik Jacobsen and Claus Thustrup Kreiner (2006b), The marginal cost of public funds: Hours of work versus labor force participation, CEPR Working Paper 5594.

Kleven, Henrik Jacobsen and Peter Birch Sørensen (2004), Labour tax reform, the good jobs and the bad jobs, Scandinavian Journal of Economics 106, 45–64.

Koskela, Erkki and Ronnie Schöb (2007), Tax progression under collective wage bargaining and individual effort determination, CESifo Working Paper 2024.

Koskela, Erkki and Jouko Vilmunen (1996), Tax progression is good for employment in popular models of trade union behaviour, Labour Economics 3, 65–80.

Lockwood, Ben and Alan Manning (1993), Wage setting and the tax system – theory and evidence for the United Kingdom, Journal of Public Economics 52, 1–29.

Mirrlees, James A. (1971), An exploration into the theory of optimal income taxation, Review of Economic Studies 38, 175–208.

Pissarides, Christopher A. (1998), The impact of employment tax cuts on unemployment and wages; the role of unemployment benefits and tax structure, European Economic Review 42, 155–183.

Ramsey, Frank P. (1927), A contribution to the theory of taxation, Economic Journal 37, 47–61.

Rutherford, Thomas F. (1998), CES preferences and technology, 89–115, GAMS Development Corporation, Washington, URL `http://www.gams.com/solvers/mpsge.pdf`.

Sørensen, Peter Birch (1997), Public finance solutions to the european unemployment problem?, Economic Policy 12, 223–251.

Sørensen, Peter Birch (1999), Optimal tax progressivity in imperfect labour markets, Labour Economics 6, 435–452.

Tuomala, Matti (1990), Optimal income taxation and redistribution, Oxford: Oxford University Press.

Appendix

A.1 Calibration

A.1.1 Income elasticity of labour supply

Originating from the homothetic CES function (2), the demand functions are homogeneous of degree one in disposable extended income. We thus have[26]

$$\varepsilon_{F_e,Y_E} = 1.$$

From this we can derive the income elasticity of labour supply. To be precise, we add a small amount of non-labour income, Y_0 to the extended income in (3)

$$Y_E = w\left[H(1 - t_L^a) + F_e(1 - t_L^m)\right] + Y_0,$$

and calculate the per cent change of labour supply with respect to a variation in Y_0 that would increase $Y = wH(1 - t_L^a) + Y_0$ by one percent, if labour supply did not react.

$$\eta_{HY} := \varepsilon_{H,Y} = \frac{\partial H}{\partial Y_0}\frac{Y}{H} = \varepsilon_{H,F_e}\varepsilon_{F,Y_e}\frac{\partial Y_E}{\partial Y_0}\frac{Y_e}{Y_E}$$

We have

$$\varepsilon_{H,F_e} = -\frac{T - H}{H}$$

$$\varepsilon_{F,Y_e} = \frac{\partial Y_E}{\partial Y_0} = 1$$

and therefore

$$\eta_{LY} = -\frac{T - H}{H}\frac{Y_e}{Y_E}$$

$$= -\frac{T - H}{H}\frac{wH(1 - t_L^a)}{w\left[H(1 - t_L^a) + (T - H)(1 - t_L^m)\right]}$$

η_{LY} is a parameter that can be observed empirically, and we use it to determine T, the (unobservable, disposable) time endowment. Solving for T, as a multiple of initial labour supply, gives

$$\frac{T}{H} = \frac{(1 - t_L^a) - \eta_{LY}(t_m - t_L^a)}{(1 - t_L^a) + \eta_{LY}(1 - t_L^m)}$$

$$= 1 - \frac{\eta_{LY}(1 - t_L^a)}{\eta_{LY}(1 - t_L^m) + (1 - t_L^a)} \tag{11}$$

---

[26]$\varepsilon_{x,y}$ is used as a shorthand for the elasticity $\frac{\partial \log x}{\partial \log y}$.

For small, negative values of η_{LY} (following Ballard (2000), I take -0.1 as the default value), $T > L$ is warranted. At the same time, small absolute values of η_{LY} will result in a small amount of disposable leisure. In the case of proportional taxes ($t_L^m = t_L^a$), (11) reduces to

$$\frac{T}{L} = 1 - \frac{\eta_{LY}}{1 + \eta_{LY}} = \frac{1}{1 + \eta_{LY}} \tag{12}$$

For $\eta_{LY} = -0.1$, we thus arrive at $T \approx 1.1$. This may seem overly little: only 4 hours of disposable leisure in relation to a standard work week of 40 hours. In ad-hoc specifications, one rather finds values between 1.5 and 2. However, this would lead to income elasticities of labour supply which are far beyond what we empirically observe. This point has forcefully been made by Ballard (2000).

A.1.2   Wage elasticity of labour supply

With the time endowment determined by the income elasticity of labour supply, we proceed with calibrating the value of the elasticity of substitution between material consumption and leisure, using the wage elasticity of labour supply (to be precise: the elasticity of labour supply with respect to the marginal after-tax wage), η_{Hw}. This is calculated as

$$\eta_{Hw} = \varepsilon_{H,\tilde{w}} = -\frac{T-H}{H}\varepsilon_{H,\tilde{w}},$$

where $\tilde{w} = w(1 - t_L^m)$. The elasticity of leisure demand with respect to the marginal after-tax wage can be decomposed into a substitution effect and an income effect:

$$\eta_{Hw} = -\frac{T-H}{H}\left[-\sigma\theta_C - (1 - \theta_C) + \frac{w(1 - t_L^m)T}{Y_E}\right]$$

Solving for σ gives

$$\sigma = \frac{\eta_{Hw} - \frac{T-H}{H}\left((1 - \theta_C) - \frac{w(1-t_L^m)T}{Y_E}\right)}{\frac{T-H}{H}\theta_C}, \tag{13}$$

which is our calibration equation. To develop a feeling for magnitudes, we again consider the special case with $t_L^m = t_L^a = t_L$. Then we have

$$\frac{w(1 - t_L)T}{Y_E} = 1$$

and (13) simplifies to

$$\sigma = \frac{\eta_{Hw} + \frac{T-H}{H}\theta_C}{\frac{T-H}{H}\theta_C} = 1 + \frac{\eta_{Hw}}{\frac{T-H}{H}\theta_C} \tag{14}$$

Further simplification of (14) is achieved by observing that in this case

$$\theta_C = \frac{H}{T},$$

which yields[27]

$$\sigma = 1 + \frac{T}{T - H}\eta_{Hw}$$

Finally, we insert (12), which leaves us with

$$\sigma = 1 - \frac{\eta_{Hw}}{\eta_{HY}}$$

This shows that the inclusion of η_{LY} in the calibration makes the outcome for σ more volatile. With an exogenous, relatively large T/H ratio, a small value of η_{Lw} would have warranted a small deviation of σ from one.[28] With η_{LY} additionally appearing in the equation, σ may easily assume much higher values. The default elasticity values, $\eta_{Lw} = 0.1$ and $\eta_{LY} = -0.1$, produce $\sigma = 2$.

A.1.3 Elasticity of participation

The distribution of the U_0's over the population must be calibrated. We have the actual participation rate and the elasticity of labour supply at the extensive margin as our empirical basis. This is sufficient to calibrate the distribution of the fixed costs locally (at the point of actual participation), but not globally. The rest of the distribution must be fixed by some functional assumption. We assume that fixed costs are uniformly distributed between U_0^- and U_0^+. For fixing the values of these bounds, we first have to calculate the change in U_l produced by an exogenous variation in the wage. We consider the case of an isolated change in the wage of the respective individual in the case of employment. In this case, the unemployment rate and the utility in case of unemployment may be considered constant. This would not be the case for a general change in the wage, which applies to all individuals. In

---

[27]This is also what you have in Rutherford (1998), if you leave out the upper nest with the consumption-savings decision (assuming that the savings ratio is zero).

[28]I follow Sørensen (1999) and assume a value of 0.1 for η_{Lw}. The meta study of Evers et al. (2005) suggests a somewhat higher elasticity, but it is difficult to distil a "core" value from this study.

terms of elasticities, we then have

$$
\begin{aligned}
\varepsilon_{U_l,w} &= \frac{(1-u)U_e}{U_l}\varepsilon_{U_e,w} = \frac{(1-u)U_e}{U_l}\left(\varepsilon_{Y_e,w} - \varepsilon_{p_U,w}\right) \\
&= \frac{(1-u)U_e}{U_l}\left(\frac{wT(1-t_L^m)}{Y_e} - \frac{wF_e(1-t_L^m)}{Y_e}\right) \\
&= \frac{(1-u)U_e}{U_l}\frac{wH(1-t_L^m)}{Y_e}
\end{aligned}
$$

The elasticity of labour supply at the extensive margin (N is the number of participating persons) is calculated as

$$
\eta_{Nw} = \varepsilon_{N,U_l}\varepsilon_{U_l,w} = h\frac{(1-u)U_e}{N}\frac{wH(1-t_L^m)}{Y_e},
$$

where h is the density of the fixed cost distribution. Solved for h, we get

$$
h = \eta_{Nw}\frac{NY_e}{(1-u)U_e wH(1-t_L^m)}.
$$

This is evaluated at the initial point, with η_{Nw} set to 0.2, following Kleven and Kreiner (2006b).[29] h is then treated as a constant in the counterfactual simulations. This means that the elasticity at the extensive margin is precisely reproduced only for the initial point; off the initial situation, it is endogenous.

The bounds of the uniform distribution for h are determined as

$$
\begin{aligned}
U_0^- &= \bar{U}_l - \frac{\bar{N}}{h} \\
U_0^+ &= \bar{U}_l + \frac{N_0 - \bar{N}}{h}
\end{aligned}
$$

where N_0 is the total population and $\bar{N}$ is initial participation. Finally, counterfactual participation is calculated as

$$
N = \bar{N} + h(U_l - \bar{U}_l)
$$

---

[29]Kleven and Kreiner (2006b, p.18-20) survey the current state of empirical evidence on the elasticity at the extensive margin. It is particularly difficult to calibrate a model with a representative agent to these elasticities, because they differ considerably by household type. The value of 0.2 is the aggregate average in Kleven and Kreiner's core scenario.

A.1.4 Nash bargaining equation

From the Nash bargaining equation (8),

$$\max_{w} \quad \Omega = (U_e - U_a)^\lambda \pi,$$

we derive the first-order condition

$$\frac{\epsilon \Omega}{\epsilon w} = \lambda \varepsilon_{U_e,w} \frac{U_e}{U_e - U_a} - \frac{wL}{\pi} = 0. \tag{15}$$

Assuming that the initial state is an equilibrium, (15) must hold and the value of λ is determined as

$$\lambda = \frac{\bar{w}\bar{L}}{\bar{\pi}\bar{\varepsilon}_{U_e,w}} \frac{\bar{U}_e - \bar{U}_a}{\bar{U}_e}.$$

Given that the income shares of labour and profits are constant in a Cobb-Douglas world, (15) boils down to the condition

$$\varepsilon_{U_e,w} \frac{U_e}{U_e - U_a} = const. \tag{16}$$

We might have started out with a utilitarian union instead, which also considers the level of employment. Then the Nash function would be

$$\Omega' = [(U_e - U_a)N]^\lambda \pi,$$

and the corresponding first-order condition

$$\varepsilon_{\Omega',w} = \lambda \left(\varepsilon_{U_e,w} \frac{U_e}{U_e - U_a} + \varepsilon_{L,w} - \varepsilon_{H,w} \right) - \frac{wL}{\pi} = 0, \tag{17}$$

where the additional terms are the elasticities of employment, L, and hours, H with respect to the wage. As long as these are constant or almost constant, we would essentially be back at (16), although λ has a different value now.

One could also try more general Nash functions with variable weights for individual utility and employment, such aslike

$$\max_{w} \quad \Omega = \left[(U_e - \bar{U})^\mu N^{(1-\mu)} \right]^\lambda \pi,$$

which includes both the utilitarian union ($\mu = 0.5$) and the insider model of (8) with $\mu = 1$ (see Graafland et al. 2001, ch. 7). This, however, would leave us with two parameters, μ and λ, which cannot be calibrated in a single first-order condition without further information.

A.2 OECD data sources

The entries in Table 1 have been generated in the following way:

- "s_L": share of labour in value added. From OECD "Annual National Accounts of OECD countries, Vol. 2", Issue 2005, Table 2: Gross domestic product: income approach. "1. Compensation of employees" / ("1. Compensation of employees" + "31. Gross operating surplus and gross mixed income"

- "t_C", "t_L^a", "t_π": effective average tax rates on consumption, labour and capital income. From OECD "Annual National Accounts of OECD countries, Vol. 2", Issue 2005, and OECD "Revenue Statistics", Issue 2004. Calculated as proposed in Mendoza et al. (1994) and further developed by Gurgel et al. (2007). In order to better fit the tax bases identified in the model of this paper, I have used the gross instead of the net capital income as basis for the capital (profit) tax. This gives substantially lower capital tax rates than those reported in the papers cited.

- "CRIP": coefficient of residual income progression. Calculated as $(1 - t_L^m)/(1 - t_L^a)$, where t_L^m and t_L^a are taken from OECD "Taxing Wages", Issue 2004, Single no child earning 100% of average production worker (APW), entries "153" and "144", respectively.

- "u": standardised unemployment rate. From OECD "Labour Force Statistics", Issue 2005, entry "(ALFS) Total labour force, All persons, Unemployment, % total labor force".

- "c": replacement rate. From OECD "Benefits and Wages", Issue 2004, Table 3.3a. (p. 102) "Average of Net Replacement Rates over 60 months of unemployment 2001, for four family types and two earnings levels, in per cent", entry "without social assistance, no children, single person".