

König, Johannes

Article — Published Version

Bias in Tax Progressivity Estimates

National Tax Journal

Provided in Cooperation with:

German Institute for Economic Research (DIW Berlin)

Suggested Citation: König, Johannes (2023) : Bias in Tax Progressivity Estimates, National Tax Journal, ISSN 1944-7477, The University of Chicago Press, Chicago, IL, Vol. 76, Iss. 2, pp. 267-289, <https://doi.org/10.1086/724186>

This Version is available at:

<https://hdl.handle.net/10419/296752>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

BIAS IN TAX PROGRESSIVITY ESTIMATES

Johannes König

Tax progressivity is central in public and political debates when questions of vertical equity are raised. Applied, structural research demands a simple way to capture it. A power function approximation delivers one parameter that captures the residual income elasticity — a summary measure of progressivity. This approximation is accurate, tractable, and interpretable, and hence immensely popular. The most common procedure to estimate this parameter, a log ordinary least squares specification, produces biased and inconsistent estimates. A nonlinear estimator solves this issue and, using different data sets, I find differences in estimates between 6 and 14 percent.

Keywords: income taxation, progressivity, nonlinear estimation

JEL Codes: H20, C51, H31

I. INTRODUCTION

The progressivity of the tax system is a central object of public and political debate when questions of vertical equity are raised. Further, it is a central characteristic of the economic environment, which determines the agents' incentives, behavior, and economic outcomes.

In life-cycle models of individual behavior, accounting for progressive taxation is essential. Over their life cycle, individuals experience unanticipated changes in the remuneration of their work, which individuals would want to insure against but cannot because insurance markets are incomplete. The government can offer additional insurance to risk-averse individuals by making the tax system progressive (insurance effect), but progressive taxation additionally induces a disincentive effect on, for example, labor supply or saving (Eaton and Rosen, 1980; Varian, 1980), which induces a trade-off. Several papers, among them Blundell, Pistaferri, and Saporta-Eksten (2016), Heathcote, Storesletten, and Violante (2014, 2017), and Kaplan

Johannes König: DIW Berlin/SOEP, Berlin, Germany (jkoenig@diw.de)

Electronically published May 19, 2023

National Tax Journal, volume 76, number 2, June 2023.

© 2023 National Tax Association. All rights reserved. Published by The University of Chicago Press on behalf of the National Tax Association. <https://doi.org/10.1086/724186>

(2012), take these two effects into account when they model individuals' life cycles. These papers introduce progressivity of the tax system by choosing a power function (Feldstein, 1969; Persson, 1983; Benabou, 2000, 2002) to approximate the relationship between gross and net income. The exponent of this power function gives the degree of progressivity of the tax system (elasticity of residual income due to Musgrave and Thin, 1948). Economists estimating structural models find this function attractive because of three features: (1) *fit* to the empirical relationship, (2) *tractability* in (structural) economic models, and (3) *interpretability* because it directly states progressivity.

The power function approximation has the following form:

$$y_{it}^{net} \approx \chi y_{it}^{1-\tau}, \quad (1)$$

where y_{it}^{net} is the net income of individual i at time t , and y_{it} is gross income. The exponent $1 - \tau$ gives the progressivity of the tax schedule. If $1 - \tau$ is smaller than one, the tax schedule is progressive; if $1 - \tau$ is larger than one, the system is regressive. Finally, if τ is equal to zero, the tax function is linear and the parameter χ gives the net-of-tax rate.¹

This paper is concerned with the estimation of $1 - \tau$. The most popular way to estimate $1 - \tau$ is to take the logarithm on both sides of Equation (1) and run ordinary least squares (Log OLS). However, the log transformation introduces a dependence between regressor and error term that generally leads to inconsistent estimates of the parameters of the power function when estimating by Log OLS — a result first shown in the context of gravity models by Silva and Tenreyro (2006). I use the Poisson pseudo maximum likelihood (PPML) estimator to estimate τ , which is nonlinear and uses data in levels.

To illustrate the empirical importance of the bias, I use three different data sets from two papers. Following Heathcote, Storesletten, and Violante (2017), I use the Panel Study of Income Dynamics (PSID) data and data released by the Congressional Budget Office (CBO) and, based on Kaas et al. (2021), I use the German Socio-Economic Panel (SOEP). The PSID is a survey data set and offers the ability to construct taxes and transfers on the household level using the National Bureau of Economic Research's (NBER) tax calculator TAXSIM (Feenberg and Coutts, 1993). Using the PSID, the estimate of $1 - \tau$ based on Log OLS is 0.819, while based on PPML it is 0.871, a 6 percent difference. The PSID, however, is not representative at the top of the income distribution, taxes are imputed through TAXSIM, and it only

¹ The power function approximation does not necessarily provide the best fit to the data. Guner, Kaygusuz, and Ventura (2014) and Kurnaz and Yip (2022) fit several parametric forms, among them the power function approximation. Both find that although it offers a good fit, it is not the best performer. For example, Guner, Kaygusuz, and Ventura (2014) show that the more complex approximation by Gouveia and Strauss (1994) provides a better fit to average and marginal income tax rates in the United States. Nevertheless, because of the attractive feature of interpretability, researchers often choose the power function approximation and will continue to do so.

provides a subset of in-kind transfers. The CBO data, which are based on tax returns, address these concerns. Here, the difference is larger: the Log OLS estimate is 0.790 and the PPML estimate is 0.904, a 14 percent difference. To broaden the perspective beyond the United States and to examine a tax system that is more progressive and redistributive than in the United States, I turn to Germany and the SOEP. Based on Log OLS, the estimate of $1 - \tau$ for the total population is 0.654, while the PPML estimate is 0.739, a 13 percent difference. However, the use of the power function approximation and Log OLS extends far beyond these three examples. Table 1 lists several prominent, but by no means all, papers in the area of taxation using the power function approximation.

While the differences in estimates of $1 - \tau$ are important in and of themselves, they become even more important when they are used for subsequent calculations, for example, when they are used to calculate tax rates or fed into a structural model. I use Heathcote, Storesletten, and Violante (2017) to illustrate this because the paper not only estimates $1 - \tau$ but also contains an optimal progressivity exercise. The authors provide a closed-form expression for social welfare, which crucially depends on the estimate of τ . Thus, the paper provides a suitable environment to illustrate how downstream results depend on the estimation of tax progressivity. The authors characterize the progressivity of the optimal income tax according to

Table 1
Selected Literature Using the Power Function Approximation

Citation	Estimates	Uses Log	Relies on
	$1 - \tau$	OLS	Estimate of $1 - \tau$
Li and Sarte (2004)	Yes	No	Yes
Kaplan (2012)	Yes	Yes	Yes
Guner, Kaygusuz, and Ventura (2014)	Yes	Yes	Yes
Heathcote, Storesletten, and Violante (2014)	Yes	Yes	Yes
Blundell, Pistaferri, and Saporta-Eksten (2016)	Yes	Yes	Yes
Heathcote, Storesletten, and Violante (2017)	Yes	Yes	Yes
García-Miralles, Guner, and Ramos (2019)	Yes	Yes	Yes
Holter, Krueger, and Stepanchuk (2019)	Yes	Yes	Yes
Heathcote, Storesletten, and Violante (2020)	Yes	Yes	Yes
Heathcote and Tsujiyama (2021)	No	Yes	Yes
Kaas et al. (2021)	Yes	Yes	Yes
Wu (2021)	Yes	Yes	Yes
Kurnaz and Yip (2022)	Yes	No	Yes

Note: A paper “relies on estimate of $1 - \tau$ ” if the authors use their estimate of $1 - \tau$ or some estimate of $1 - \tau$ to make further calculations, for example, to calculate tax burdens or in an optimal progressivity exercise.

a utilitarian objective. They find that progressivity is too high in the status quo and that switching to the optimum would bring about a welfare gain. Using the PPML estimate of $1 - \tau$, this welfare gain is about 67 percent smaller.

The rest of the paper is organized as follows. In Section II, I formally introduce the power function approximation and outline the estimation framework. Section III introduces the three data sets for estimation. Section IV presents the estimation results, and Section V illustrates the importance of the findings. Section VI concludes.

II. THE POWER FUNCTION APPROXIMATION

I compare two techniques to estimate $1 - \tau$: Log OLS and PPML. To see why the two techniques will deliver different estimates of progressivity and different empirical fits, consider the stochastic version of Equation (1):²

$$\begin{aligned} y_{it}^{net} &= \chi y_{it}^{1-\tau} \eta_{it} \\ \eta_{it} &= 1 + \epsilon_{it}, \end{aligned} \tag{2}$$

where ϵ_{it} is an error term, which I will define later. The above formulation shows that a multiplicative error formulation can be translated into an additive error formulation and vice versa.

The assumption necessary for the identification of Equation (2) is that the conditional mean of y_{it}^{net} equals $\chi y_{it}^{1-\tau}$. This appears to be a reasonable and natural assumption in this setting, because it expresses in precise, statistical terms what Equation (1) states.³ Papers that use the power function as an approximation of y_{it}^{net} appear to be in tacit agreement with this assumption about the conditional mean of y_{it}^{net} .⁴ Restating this in terms of a requirement on the error term, it should hold that $E[\eta_{it}|y_{it}] = 1$.

Under this assumption, it is possible to identify the parameters of Equation (2) using a nonlinear estimator. However, the consistent estimation of the parameters of Equation (2) is possible by Log OLS only when ϵ_{it} has *one* specific form. When $\epsilon_{it} = \nu_{it}$, with ν statistically independent of $\log y$, Log OLS identifies the parameters of the power function.

$$\begin{aligned} \log y_{it}^{net} &= \log \chi + (1 - \tau) \log y_{it} + \log(1 + \epsilon_{it}) \\ &= \log \chi + (1 - \tau) \log y_{it} + \log(1 + \nu_{it}). \end{aligned} \tag{3}$$

² This exposition follows Silva and Tenreyro (2006).

³ In fact, this is precisely the statement that Silva and Tenreyro (2006) make about economic models, when they point out that: "All that can be expected is that they hold on average" (p. 643). This notion is precisely captured by the requirement that the conditional mean of y_{it}^{net} equals $\chi y_{it}^{1-\tau}$.

⁴ Assumptions about the error are seldom stated in the literature.

But this implies that in levels

$$y_{it}^{net} = \chi y_{it}^{1-\tau} + \chi y_{it}^{1-\tau} \nu_{it}. \quad (4)$$

This composite error term implies a specific type of heteroskedasticity, where the conditional variance of y_{it}^{net} is proportional to $(\chi y_{it}^{1-\tau})^2$. Whether the heteroskedasticity of the error term satisfies this specific functional form is ultimately an empirical question because the joint distribution of net and gross incomes depends on the tax and transfer schedule, deductions, tax credits, and individual behavior. However, the variance of net incomes should generally rise with gross income. At the lower end of the distribution, the variation in tax burdens — for example, due to child tax credits or similar provisions — is fairly limited. In the upper part of the distribution, it is expected that there is wide dispersion in net incomes depending on the ability to use deductions or tax credits. But there is no reason to believe that it should be proportional to $(\chi y_{it}^{1-\tau})^2$. In Section IV, tests of the null hypothesis that the error term has this specific form of heteroskedasticity reject it at conventional levels of statistical significance. But in some cases, the use of Log OLS may be innocuous. I show that when I exclude the upper half of the gross income distribution from the PSID data, the Log OLS and PPML estimates converge.

To see why Log OLS generally does not consistently estimate the parameters of Equation (2), consider the case of an additive error in levels, such that $\epsilon_{it} = \nu_{it}/(\chi y_{it}^{1-\tau})$. Then after log transformation

$$\log y_{it}^{net} = \log \chi + (1 - \tau) \log y_{it} + \log \left(1 + \frac{\nu_{it}}{\chi y_{it}^{1-\tau}} \right).$$

Clearly, the log transformation introduces a dependence between the error term and the regressor, and one cannot obtain a consistent estimate of $(1 - \tau)$ using Log OLS.

The parameters of Equation (2) can be consistently estimated with a nonlinear pseudo maximum likelihood estimator. It solves

$$\sum_{j=1}^J (y_j^{net} - \hat{\chi} y_j^{1-\tau}) y_j = 0, \quad (5)$$

where j is an index that spans all observations. This estimator happens to be numerically equal to the PPML. Note that, although Poisson regression is usually used to model count data, this is not assumed here and not necessary for the consistent estimation of the parameters. All that is required for consistent estimation is the correct specification of the conditional mean. Thus, y^{net} does not have to be count data and does not need to follow a Poisson distribution (Gourieroux, Monfort, and Trognon, 1984; Wooldridge, 2010, Ch. 18).⁵

⁵ However, the standard errors produced by common Poisson regression packages will not be correct because they are computed under the assumption that the conditional variance equals the conditional

In practice, the estimates can be computed by running a Poisson regression of y_{it}^{net} on $\ln y_{it}$ in statistical software like Stata or R.⁶ Silva and Tenreyro (2006) propose this estimator for these types of data and test it in Monte-Carlo experiments with different types of heteroskedasticity. They show that it performs well in many different heteroskedasticity settings, including the one where Log OLS is consistent. Other nonlinear estimators — like nonlinear least squares, for example — are consistent but vary in performance depending on the type of heteroskedasticity prevalent in the data.⁷

III. DATA

In my empirical exercises, I employ three data sets: (1) the data based on the PSID in Heathcote, Storesletten, and Violante (2017), (2) my construction of the CBO data set on net and gross incomes following Heathcote, Storesletten, and Violante (2017), and (3) data from the SOEP constructed using the replication files of Kaas et al. (2021). Heathcote, Storesletten, and Violante (2017) use the tax calculator TAXSIM (Feenberg and Coutts, 1993) provided by NBER to calculate net incomes. The CBO data set is based on the Statistics of Income and therefore is based on a nationally representative sample of income tax returns and contains incomes and tax liabilities. The SOEP data set (Goebel et al., 2019; Schröder et al., 2020) contains gross and net incomes based on a tax calculator maintained at the German Institute for Economic Research (Grabka, 2020). For more details on the data, see Appendix A (Appendixes A–F are available online).

IV. RESULTS

A. PSID Data

Table 2 presents the Log OLS and PPML estimates of $1 - \tau$, bootstrapped standard errors, as well as fit statistics. The estimate of $1 - \tau$ using Log OLS is smaller (0.819), implying a more progressive tax system, than the one obtained by PPML (0.871).

Goodness of fit (GOF) is assessed by computing errors from observed data and model predictions, that is, $e_{it} = y_{it}^{net} - \hat{y}_{it}^{net}$, where \hat{y}_{it}^{net} is the respective prediction of net income.⁸ I then compute the root mean squared error (RMSE) and the mean absolute error (MAE). The RMSE indicates that the model estimated using PPML improves on the fit of Log OLS by about 31 percent. The MAE derived from PPML is also smaller, but the improvement is more modest at roughly 9 percent.

mean. Wooldridge (2010, Ch. 18) shows how to compute appropriate standard errors. In Stata these can be produced with the option “robust” and in R using the library “sandwich.”

⁶ In Stata the command is “poisson” and in R it is “glm(.,family = “poisson”).”

⁷ Not all papers use Log OLS to estimate $1 - \tau$. For example, Kumaz and Yip (2022) use nonlinear least squares.

⁸ In the case of Log OLS, I adjust for Jensen’s inequality in the prediction. See Rainey (2017).

Table 2
Progressivity Estimates and GOF in Heathcote,
Storesletten, and Violante (2017) Data Set

	Log OLS	PPML
$1 - \tau$	0.819 (0.00493)	0.871 (0.00852)
RMSE	17,878	12,332
MAE	4,457	4,061
Obs.	12,875	12,875

Note: This is my own calculation based on the Heathcote, Storesletten, and Violante (2017) replication data set. Bootstrapped standard errors in parentheses are based on 500 replicates. Pooled observations are from 2000 to 2006.

To put these differences in perspective, I calculate the average income-weighted marginal tax rate (AMTR; Barro and Sahasakul, 1983), as Heathcote, Storesletten, and Violante (2017) do. They find an AMTR of about 34 percent and report that Barro and Redlick (2011) find a slightly higher AMTR of about 37 percent. Calculating the AMTR using the PPML estimates, I match the AMTR in Barro and Redlick (2011) exactly. I find a slightly higher AMTR because, although the estimate of $1 - \tau$ is larger, the estimate of χ is lower compared with Heathcote, Storesletten, and Violante (2017).

I show the absolute errors $|e_{it}|$ produced by either procedure along the distribution of gross income in Figure 1. The dark gray diamonds indicate errors produced by Log OLS; the light gray circles indicate errors for PPML. For large values of gross income Log OLS produces much larger errors than PPML, while for small values of gross income the difference is small. With lower estimated progressivity PPML appears to improve the fit for larger net incomes. This is intuitive because past the final tax bracket, the tax system becomes basically linear for high-income earners.

A boxplot of net incomes over gross income deciles in Figure 2 helps to understand the pattern of heteroskedasticity in net incomes. The figure shows that the variance of net incomes grows over the deciles, with the largest amount of variation in the 10th decile.

In Appendix B, I show several robustness exercises. I successively restrict the sample, dropping more and more high gross income observations from the data set. When 50 percent of the sample has been dropped, the two estimates converge. Thus, even though the results become similar for lower levels of income, generally, Log OLS still tends to estimate smaller values of $1 - \tau$. Further, restricting to the population with only positive taxes also leads to a convergence in estimates.⁹ In this setting,

⁹ García-Miralles, Guner, and Ramos (2019) use the power function to model the tax system, but only for positive taxes. Thus, it is interesting to compare estimates on this restricted sample.

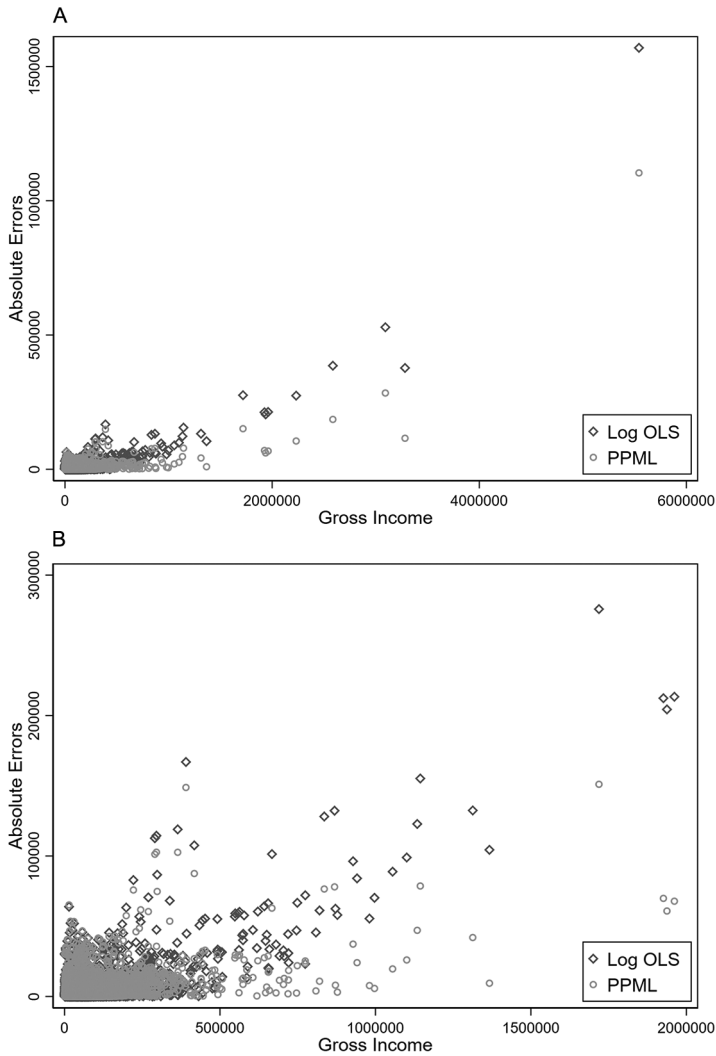


Figure 1. Absolute errors along the distribution of gross income. This is my own calculation based on the Heathcote, Storesletten, and Violante (2017) replication data set. I plot the absolute errors, that is, the absolute difference between observed and predicted net income from Log OLS and PPML. Panel A shows the errors for the full data, while panel B shows the errors in the range up to \$2 million of gross income.

Log OLS gives an estimate of $1 - \tau$ that is larger and statistically significantly distinct from the PPML estimate. While the RMSE of Log OLS is somewhat smaller than the RMSE of PPML, the MAE of PPML is still smaller than the one of Log OLS. This is likely due to the fact that Log OLS now strongly reacts to the large errors at the top of the gross income distribution. Thus, although the two estimates converge

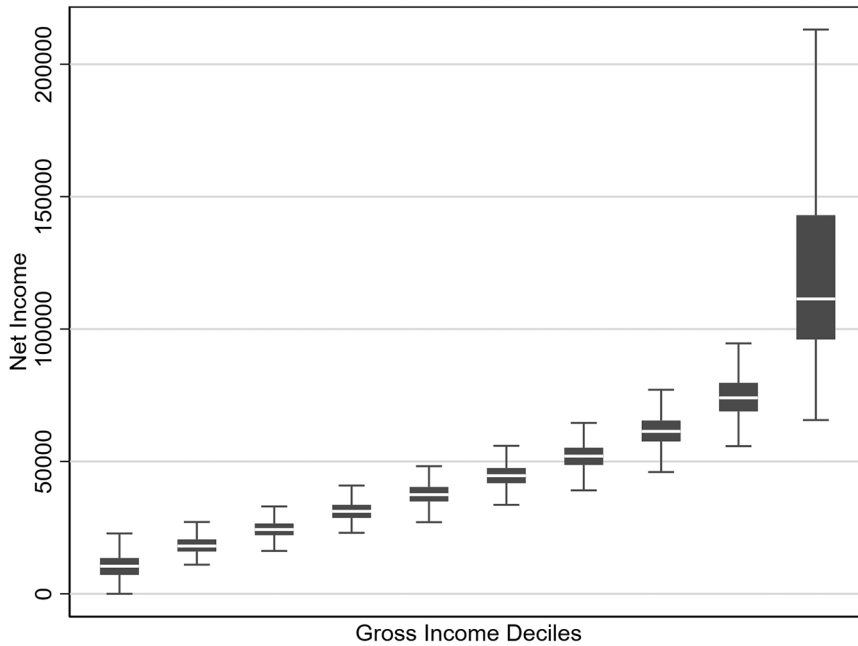


Figure 2. Boxplot of net incomes over deciles of gross income using Heathcote, Storesletten, and Violante (2017) data set. This is my own calculation based on the Heathcote, Storesletten, and Violante (2017) replication data set. Figures are in nominal dollars.

as the sample is restricted to larger gross incomes, the fundamental issue, that is, the bias due to the pattern of heteroskedasticity, cannot be fully alleviated in this manner. Further, by restricting the range in this way, the interpretation of $1 - \tau$ becomes less clear. Finally, year-by-year estimates confirm the patterns of Table 2: Log OLS produces consistently smaller estimates of $1 - \tau$ compared with PPML. Further, PPML delivers a better fit in terms of both the RMSE and the MAE in all years.

1. A Park-Type Regression Test of the Log Transformation

I have shown that PPML produces significantly different progressivity estimates compared with Log OLS and that the fit of PPML is consistently better than the fit of Log OLS. To assess whether the log transformation is appropriate and whether Log OLS will consistently estimate $1 - \tau$, I perform a Park-type regression test (Park, 1966). *If and only if* the conditional variance of y_{it}^{net} is proportional to $(\chi y_{it}^{1-\tau})^2$, Log OLS will be consistent. I run the auxiliary regression:

$$\log (y_{it}^{net} - \hat{y}_{it}^{net})^2 = \alpha + \beta \log \hat{y}_{it}^{net} + u_{it}, \tag{6}$$

Table 3
Park-Type Regression Test

β	0.738 (0.02668)
F -test: $\beta = 2$	2,237.94
p value	0.00
Obs.	12,875

Note: This is my own calculation based on the Heathcote, Storesletten, and Violante (2017) replication data set. Robust standard errors are in parentheses. Pooled observations are from 2000 to 2006.

where \hat{y}_{it}^{net} is the predicted value of y_{it}^{net} . *If and only if* β is equal to 2, the condition for Log OLS consistently estimating $1 - \tau$ is met. I use PPML to predict \hat{y}_{it}^{net} and then estimate Equation (6) by OLS. Using an F -test, I evaluate the null hypothesis of $\beta = 2$. Table 3 shows the estimated β is equal to 0.738 and the F -test rejects the null hypothesis at conventional levels of significance.

2. Nonparametric Estimation of Progressivity

The progressivity parameter can also be estimated directly and nonparametrically. Using marginal tax rates, one can compute the residual income elasticity — the equivalent of $1 - \tau$ — for each observation in the PSID data set (Musgrave and Thin, 1948; Jakobsson, 1976). As the name suggests, the residual income elasticity measures the percentage change in net income for a given percentage change in gross income. The formula for the residual income elasticity for each observation is as follows:

$$rie_{it} = (1 - mtr_{it}) \frac{y_{it}}{y_{it}^{net}}, \quad (7)$$

where rie_{it} is the residual income elasticity and mtr_{it} is the marginal tax rate.

The replication data set by Heathcote, Storesletten, and Violante (2017) does not contain marginal tax rates. Thus, I replicate the data set using PSID data and TAXSIM. Appendix C gives the details on my replication of the Heathcote, Storesletten, and Violante (2017) data set. TAXSIM computes marginal tax rates with the method of finite differences with very slight variation in the taxable amounts, usually 1 cent.¹⁰ TAXSIM reports the marginal federal and marginal state tax rates. I add both, which gives the total marginal tax rate corresponding to mtr_{it} . The replication data set has

¹⁰ See <https://taxsim.nber.org> for a more in-depth documentation.

12,886 observations for which rie_{it} could be calculated and which could be matched with an observation in the Heathcote, Storesletten, and Violante (2017) data set.

The mean of rie_{it} for the whole sample is 0.897. I plot the mean of rie_{it} for 100 quantile groups of gross income and the overall mean in Figure 3. The 10th quantile group has mean gross income of about \$20,000, while the mean in 100th quantile is about \$712,000.

The figure shows that there is a lot of heterogeneity in the residual income elasticity along the distribution of gross income, but some features emerge. First, the mean of rie_{it} is around 0.85 in the bottom of the distribution and slowly rises up to more than 0.9 in the middle of the distribution. Past the 50th quantile group, rie_{it} falls to a mean very close to 0.9. Second, although there is a lot of variability overall, rie_{it} exhibits a lot of variation for the lower quantile groups. In the middle and upper end of the distribution of gross incomes variation becomes more systematic. While it is certainly true that the tax system is more progressive at the bottom of the distribution, both in the middle, at the top, and overall, progressivity is far smaller than the Log OLS estimates suggest. Yet, even at the bottom of the income distribution, rie_{it} is not often close to the Log OLS estimate of $1 - \tau$. Overall, the progressivity estimate from PPML is closer to the mean of rie_{it} .

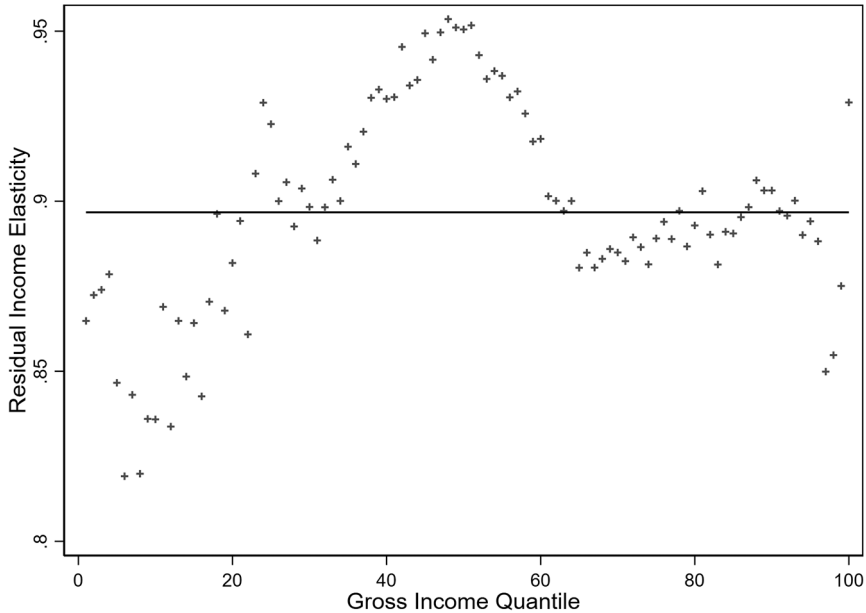


Figure 3. Residual income elasticity along the distribution of gross income. This is my own calculation based on the PSID and TAXSIM. I plot the mean of rie_{it} for 100 quantiles of gross income. The horizontal line gives the mean, which is equal to 0.897.

Table 4
Progressivity Estimates Using CBO Data Set

	Log OLS	PPML
$1 - \tau$	0.790	0.904
Obs.	56	56

Note: This is my own calculation based on CBO (2020). Data are from 2000 to 2006.

B. CBO Data

While the PSID is a benchmark data set, it has several limitations: (1) It is not representative at the very top of the income distribution. (2) It only provides information on a subset of in-kind transfers. (3) Taxes are calculated through TAXSIM. The CBO data address all of these concerns. The CBO data are based on a nationally representative sample from the Statistics of Income data and thus come from administrative tax records. Hence, the data are representative at the top, cover in-kind transfers more completely, and taxes are directly taken from the Internal Revenue Service. The data are published in the form of year-specific tables, which report the means of gross and net incomes for several percentile groups ordered by gross income. These are the bottom four quintiles, and the percentile groups 81–90, 91–95, 96–99, and the top 1 percent. I use data from 2000 to 2006. I use the “income before transfers and taxes” measure of pregovernment incomes provided by the CBO. Table 4 presents the estimates of the progressivity parameter for a Log OLS specification and a PPML specification. I do not report standard errors because they cannot help in the interpretation of the results with this data set.¹¹

The Log OLS estimate I obtain for $1 - \tau$ is fairly close to the one reported in Heathcote, Storesletten, and Violante (2017), namely 0.790, so that the estimate τ is 0.210. The progressivity estimate using PPML is a bit higher than the analogue in the PSID at 0.904. The difference in estimates is roughly 14 percent. By more accurately capturing top earners in the data set and by including more in-kind transfers, variation in net incomes along gross incomes changes compared with the PSID data set. This additional variation at both the top and the bottom leads to a stronger divergence in Log OLS and PPML estimates. This is most likely due to the fact that the conditional variance now further diverges from the case, which would make Log OLS consistent.

The CBO data cover both the top and the bottom of the distribution more accurately. Thus, ex ante, it is not clear whether we should expect a more or a less progressive estimate. Better coverage of top earners should lead to less progressivity, but the inclusion of more in-kind transfers should lead to more progressivity. To

¹¹ This is because the data are means within certain sections of the income distribution making the data grouped.

shed more light on which of these forces is the more important one, I compare some distributional statistics in the PSID and the CBO data. I calculate the average gross income of the top quintile of households in the PSID sample and in the CBO data and find a value of about \$164,000 for the PSID and about \$250,000 for the CBO data, a 52 percent difference.¹² When I calculate these figures for the middle quintile they are about \$50,000 for the PSID and \$67,000 for the CBO, thus showing a difference of 34 percent. Hence, I conclude that the more accurate representation of top earners is the quantitatively more important force and, thus, Log OLS gives a counterintuitive result.

C. SOEP Data

The previous two data sets come from the United States, a country with high gross but also high net income inequality (Piketty and Saez, 2003; Piketty, Saez, and Zucman, 2017). It is reasonable to wonder about the extent to which the Log OLS bias will play a role in other institutional settings. Therefore, I conduct further analysis of the bias using German data. Germany is a country with both smaller gross and net income inequality and a larger extent of redistribution compared with the United States (Bartels and Waldenström, 2022). It stands to reason that with a larger degree of redistribution the shape of the conditional variance of net incomes is meaningfully different from the US case, which will influence the extent of the bias. I choose to replicate Kaas et al. (2021), who model the power function as age-dependent but also report a Log OLS estimate of $1 - \tau$ for their entire sample, which is 0.654. Their estimate supports the notion that the German tax system is far more progressive than the US tax system.

In Table 5, I provide my estimates of $1 - \tau$ and the goodness of fit measures based on the data set derived from the replication files of Kaas et al. (2021). I match the estimate of $1 - \tau$ reported by Kaas et al. (2021) exactly when I apply Log OLS. By applying PPML, the estimate of $1 - \tau$ rises by about 13 percent to 0.739. Also like the PSID, both the RMSE (about 22 percent less) and the MAE (about 18 percent less) are lower for PPML. Further, I report the value of the Park-type regression test computed from the data and find that the null hypothesis of the type of heteroskedasticity required for Log OLS to be consistent is rejected at conventional levels of statistical significance. Thus, although the SOEP data are very similar to the PSID data in terms of the data quality, income concepts, and the sample restrictions, I find that the divergence between Log OLS and PPML estimates is larger. It is reasonable to attribute this larger divergence to the more redistributive character of the German tax and transfer system as it is likely to further alter the conditional variance of net incomes.

¹² I calculate these figures by computing the average gross income for the top 20 percent in each year and then averaging that estimate over the years. I can only compare averages because the CBO only reports averages.

Table 5
Progressivity Estimates and GOF in
Kaas et al. (2021) Data Set

	Log OLS	PPML
$1 - \tau$	0.654 (0.00380)	0.739 (0.00216)
RMSE	8,732	6,851
MAE	5,485	4,515
Park-test value	41,311.67	
p value	0.00	
Obs.	112,467	112,467

Note: This is my own calculation based on the Kaas et al. (2021) replication data set. Bootstrapped standard errors in parentheses are based on 500 replicates. Pooled observations are from 1995 to 2014. Park-test value refers to the F -statistic based on the auxiliary regression in Equation (6).

Figure 4 shows the boxplots of the distribution of net income in the SOEP for the deciles of gross income. The qualitative pattern of heteroskedasticity is similar to Figure 2: the variance of net incomes rises along the deciles, with the largest variance in the 10th decile, although there is also larger variation in the bottom of the distribution.

V. RELEVANCE OF THE BIAS

A. Net Income Prediction and Tax Schedule

While the power function approximation and especially the estimate of $1 - \tau$ have diverse applications, the basic use is still as a function to generate net incomes from gross incomes. Thus, as the most straightforward way to show the relevance of the differences between Log OLS and PPML estimates, I compute net incomes as well as average and marginal tax rates.

For net incomes, I choose the set of estimates from the PSID and predict using a small synthetic data set of gross incomes shown in Table D4. The synthetic distribution contains pregovernment incomes from \$5,000 to \$3 million.¹³ Figure 5 shows the functions tracing predicted net incomes for the full distribution (Panel A) and gross incomes up to \$300,000 (Panel B).

Panel A of Figure 5 shows large differences in net incomes around \$500,000 of gross income using the PPML estimates. From there, due to differences in the exponents, the differences in net incomes get larger and larger. The figure also mirrors

¹³ The lower end roughly coincides with the sample restriction on low incomes in Heathcote, Storesletten, and Violante (2017).

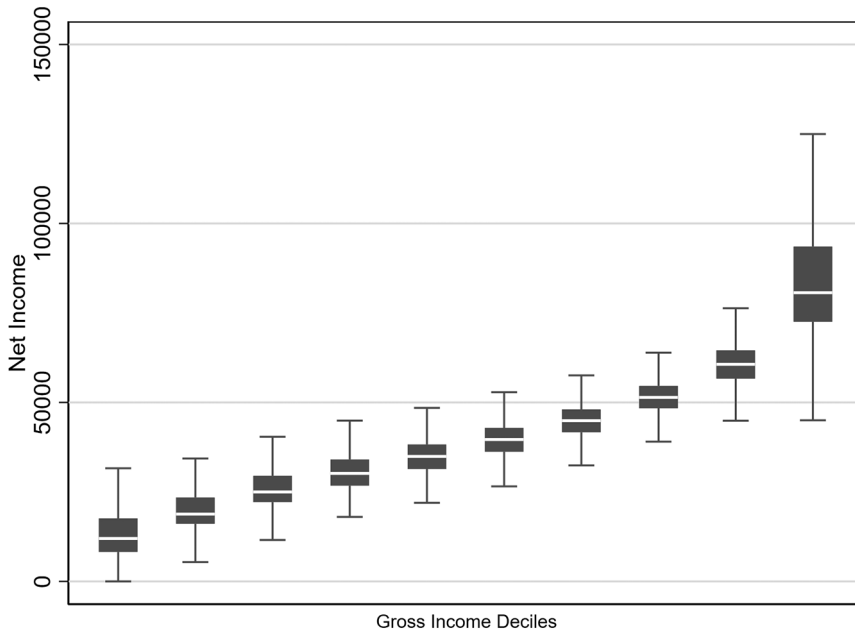


Figure 4. Boxplot of net incomes over deciles of gross income using the Kaas et al. (2021) data set. This is my own calculation based on the Kaas et al. (2021) replication data set. Figures are in 2006 euros.

the findings from Figure 1: while the fit at the bottom appears virtually the same for both functions, the fit at the top is very different.

To check whether the fit actually is different at the bottom of the distribution, I present Panel B of the figure. Here one can see that the Log OLS prediction actually starts out more progressive than the PPML prediction. This holds up until slightly above \$100,000 of income, when the PPML prediction crosses the Log OLS prediction line. But as is evident from Figure 1 and from Panel B, the differences in predictions are small and do not outweigh the increasingly worse fit that Log OLS shows at the top end of the distribution.

To check differences in the implied tax schedules, I compute average and marginal tax rates along the distribution of gross income. Average tax rates are given by $T(y_{it})/y_{it} = 1 - \chi y_{it}^{-\tau}$, while marginal tax rates are given by $\partial T(y_{it})/\partial y_{it} = 1 - \chi(1 - \tau)y_{it}^{-\tau}$. I show both, once computed with Log OLS and once with PPML estimates, in Figure 6.

Figure 6 shows the average and marginal tax rates computed with Log OLS and PPML. Both the average and the marginal tax rates are lower at the bottom and higher at the top of the distribution when using Log OLS estimates. For example, at \$300,000, Log OLS gives an average tax rate of 41.6 percent and a marginal tax rate of 52.1 percent, while PPML gives rates of 37.6 and 45.6 percent, respectively.

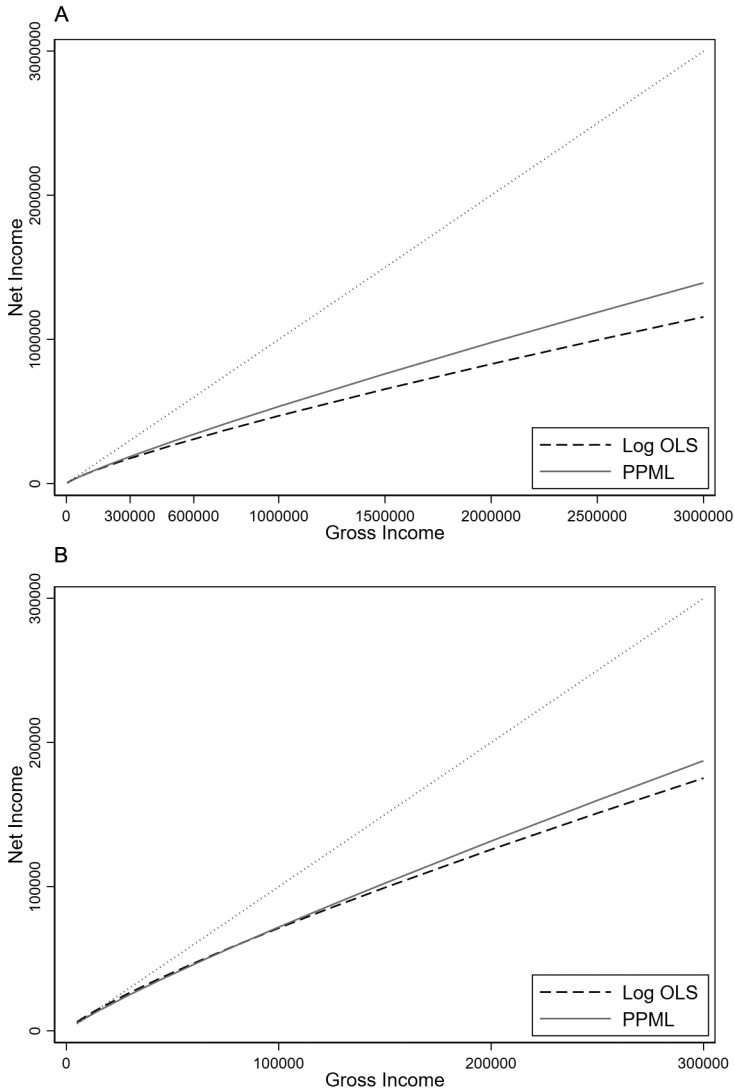


Figure 5. Predicted net incomes: (A) full distribution and (B) gross incomes up to \$300,000. This is my own calculation based on the synthetic data set in Table D4. The dotted line is the 45-degree line.

Summing up, this exercise shows that although the estimates from Log OLS and PPML deliver fairly similar results for net income and tax rates at the bottom of the distribution, once higher incomes are considered, predictions diverge considerably. Thus, for accurate net income predictions, the point in the distribution we focus on matters.

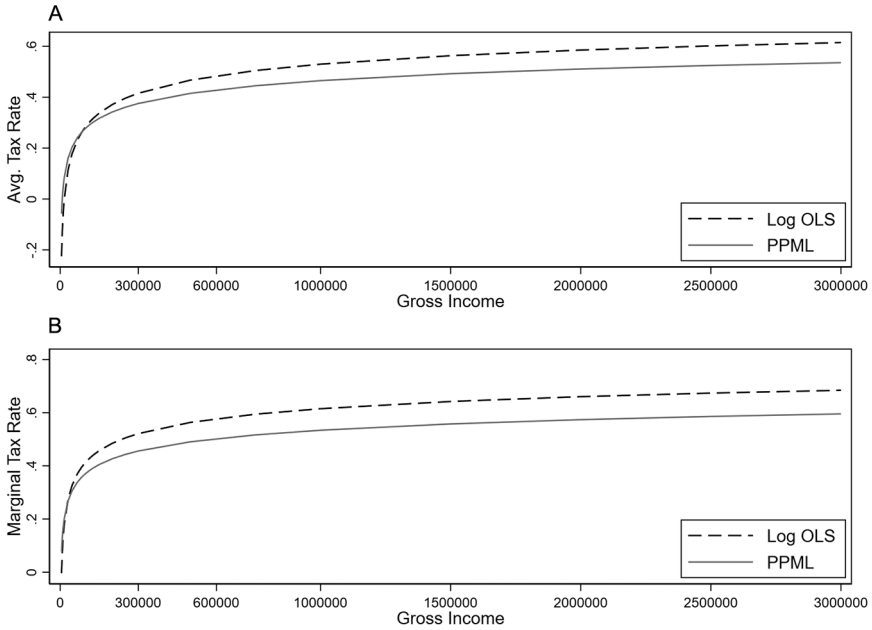


Figure 6. Predicted (A) average and (B) marginal tax rates. This is my own calculation based on the synthetic data set in Table D4.

B. Frisch Elasticity Estimates

While net income and tax rate prediction are the most straightforward uses of the power function approximation, other uses, especially for the estimate of $1 - \tau$, exist. The progressivity parameter plays a major role in life-cycle models of labor supply. In these models, it influences how shocks — from the wage process or other sources — transfer to income, consumption, and ultimately lifetime utility. Further, tax progressivity influences the labor supply reaction of individuals by attenuating the incentive to respond to shocks. Accordingly, the estimated labor supply reaction differs when one accounts for the progressivity parameter and thereby tax progressivity. Thus, we can expect a more sizable impact of the divergence in estimates between Log OLS and PPML when the two sets of estimates are used for the calculation of labor supply elasticities.

Consider the Frisch elasticity of labor supply, that is, the elasticity determining the reaction to a transitory wage shock. Estimates of the Frisch elasticity can be obtained from the estimation of an intertemporal labor supply equation. An example of such an equation is

$$\Delta \ln h_t \approx \frac{1}{\gamma + \tau} [\text{cons}_t + (1 - \tau)\Delta \ln w_t - \zeta \Delta \bar{E}_t + \Delta v_t + \eta_t], \quad (8)$$

Table 6
Unadjusted Frisch Elasticities

Tax-adj. Frisch	τ	
	Log OLS: 0.181	PPML: 0.129
	Unadj. Frisch	Unadj. Frisch
0.5	0.686	0.620
0.4	0.536	0.488
0.3	0.392	0.360
0.2	0.255	0.237

Note: This is my own calculation.

where Δ is the first difference operator, h are hours worked, $1/\gamma$ is the Frisch elasticity, $cons_t$ contains time-specific constants, w is the wage, Ξ_t are variables determining taste for work, v are shocks to the taste for work, and η is the approximation error of the marginal utility of wealth. Finally, $1 - \tau$ is the exponent of the power function approximation of the tax system.¹⁴

The reaction in $\Delta \ln h_t$ to a change in $\Delta \ln w_t$ is $(1 - \tau)/(\gamma + \tau)$. In the case without progressive taxation, when τ equals zero, the reaction would be $1/\gamma$, that is, the ordinary Frisch elasticity. However, the econometrician estimates the tax-adjusted Frisch elasticity when estimating Equation (8). Calculating the unadjusted Frisch elasticity requires an estimate of τ . In Table 6 I give an illustration of what the unadjusted Frisch elasticity would look like, for different tax-adjusted Frisch elasticity estimates and two different estimates of τ . The range of estimates of the tax-adjusted Frisch elasticities loosely follows the range of estimates presented in Keane (2011). The two estimates of τ correspond to the one estimated by Log OLS (0.181) and the one estimated by PPML (0.129)

The table shows that at low tax-adjusted Frisch elasticities, the differences in unadjusted elasticities are quite small, but they grow quickly. Heathcote, Storesletten, and Violante (2017) use an unadjusted Frisch elasticity of 0.5 in their optimal progressivity exercise, so that the implied tax-adjusted Frisch elasticities differ by about 9 percent.

Thus, depending on how the estimate of $1 - \tau$ is used in conjunction with other parameters, we can expect large differences in relevant economic preference parameters like the Frisch elasticity of labor supply. If these are used in further calculations, errors will propagate and potentially lead to incorrect quantitative or qualitative conclusions.

¹⁴ The equation can be derived from the first-order conditions of the model in Appendix E and then approximating the intertemporal Euler equation with a first-order Taylor expansion.

C. Optimal Progressivity Exercises

One of the most important uses of the estimate of $1 - \tau$ is in optimal progressivity exercises. Here, a social planner weighs the beneficial insurance effect of progressivity against the adverse incentive effects on, for example, labor supply and human capital investment. In such optimal progressivity exercises, many relevant quantities in the model, especially social welfare, directly depend on the estimate of $1 - \tau$. Thus, the relevance of the Log OLS bias is fully revealed when I conduct the optimal progressivity exercise in Heathcote, Storesletten, and Violante (2017) once with the Log OLS estimate of τ and once with the PPML estimate and compare the implications for the change in social welfare.

Heathcote, Storesletten, and Violante (2017) build a model of life-cycle labor supply and human capital accumulation and provide a closed-form expression for social welfare, which crucially depends on the level of τ .¹⁵

Heathcote, Storesletten, and Violante (2017) find that progressivity is too high compared with the utilitarian optimum. The welfare-maximizing progressivity parameter in their model is $\tau^* = 0.084$, while it is 0.181 in the status quo. The average welfare gain from switching to τ^* is equivalent to 0.63 percent of lifetime consumption. When I change the structural parameters to my estimates from Appendix F, $\tau^* = 0.071$, thus making the optimal tax system even less progressive. When I calculate the welfare gain with the PPML estimate of τ , setting $\tau = 0.129$ in the status quo, the welfare gain is only 0.21 percent of lifetime consumption. Thus, the welfare gain is about 67 percent smaller compared with the result shown in Heathcote, Storesletten, and Violante (2017). The change is so large because social welfare is a nonlinear function of τ .

Even though the qualitative conclusion in Heathcote, Storesletten, and Violante (2017) that progressivity should be lowered is not overturned, I do find that the progressivity in the status quo is much closer to the optimum than the original finding suggests. Optimal progressivity exercises and their conclusions are highly sensitive to the magnitude of the progressivity parameter.

VI. CONCLUSION

In this paper, I apply the PPML estimator of Silva and Tenreyro (2006) to estimate the power function progressivity parameter of the income tax systems in the

¹⁵ The authors also provide detailed computer programs, making it easy to reproduce their results. I adapt the MATLAB script “MainCode_HSV_QJE.m” from the replication files of Heathcote, Storesletten, and Violante (2017) for the calculations.

The results in the optimal progressivity exercise also depend on a method of moments estimation of several structural parameters from the PSID and the Consumer Expenditure Survey (CEX). They estimate variances for the heterogeneity in leisure preference and for insurable and uninsurable wage shocks, as well as the elasticity of substitution between skill types from these moments. These estimates, however, depend on the estimate of τ . Appendix F shows these parameters using $\tau = 0.129$, as estimated by PPML. I use these adjusted estimates of the structural parameters when assessing welfare gains from optimal progressivity.

United States and Germany. I compare the results with the commonly applied Log OLS procedure. The Log OLS procedure is generally biased except in a very specific case of heteroskedasticity, which can be empirically tested. Thus, whether Log OLS is biased and whether the bias matters are empirical questions.

My findings are: (1) Log OLS estimates of the progressivity parameter $1 - \tau$ are lower than the estimates obtained by PPML, but the extent depends on the data set. The change found in the PSID data (6 percent) is smaller than the change found in the CBO data (14 percent) and the SOEP data (13 percent). (2) The Park-type regression test rejects the null hypothesis that the specific case of heteroskedasticity, which would make Log OLS consistent, is satisfied. (3) The Log OLS predictions of net income produce larger RMSEs and larger MAEs when using PSID and SOEP data. (4) The evidence from the SOEP suggests that when the tax system is more progressive, the extent of the bias is larger.

Section V shows that the differences between Log OLS and PPML estimates are economically meaningful. For instance, the welfare gains shown in the optimal progressivity exercise of Heathcote, Storesletten, and Violante (2017) are 67 percent smaller using the PPML estimates.

Researchers who want to apply the power function approximation should estimate it using PPML, as there is no downside. PPML consistently estimates the progressivity parameter, even in the heteroskedasticity scenario where Log OLS is consistent. Further, as Silva and Tenreyro (2006) and I point out, running a Poisson regression of net income on the log of gross income delivers the PPML estimates, making it as easy as running Log OLS. Finally, researchers can test whether the use of Log OLS is innocuous with the Park-type regression test.

ACKNOWLEDGMENTS

I want to express my gratitude to Leo Kaas and Georgi Kocharkov for providing me with both the codes and data files from Kaas et al. (2021) to exactly replicate their findings. For helpful comments and discussions, I thank Thomas Aronsson, Richard Blundell, Patrick Burauel, Giacomo Corneo, Giulio Fella, Michael Graber, Katharina Jenderny, Robin Jessen, Loukas Karabarbounis, Musab Kurnaz, Gauthier Lanot, Maximilian Longmuir, Maria Metzger, Sebastian Schmitz, Carsten Schröder, Michael Stepner, and Alexandros Theloudis as well as participants at the Annual Congress of the International Institute of Public Finance 2019, the Annual Congress of the National Tax Association 2019, and the Annual Congress of the European Economic Association 2020.

DISCLOSURES

The author has no relevant or material financial interests that relate to the research described in this paper.

REFERENCES

- Barro, Robert J., and Charles J. Redlick. 2011. "Macroeconomic Effects from Government Purchases and Taxes." *Quarterly Journal of Economics* 126 (1), 51–102.
- Barro, Robert J., and Chaipat Sahasakul. 1983. "Measuring the Average Marginal Tax Rate from the Individual Income Tax." *Journal of Business* 56 (4), 419–452.
- Bartels, Charlotte, and Daniel Waldenström. 2022. "Inequality and Top Incomes." In Zimmermann, Klaus F. (ed.), *Handbook of Labor, Human Resources and Population Economics*, 1–28. Springer, Cham.
- Benabou, Roland. 2000. "Unequal Societies: Income Distribution and the Social Contract." *American Economic Review* 90 (1), 96–129.
- Benabou, Roland. 2002. "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica* 70 (2), 481–517.
- Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. "Consumption Inequality and Family Labor Supply." *American Economic Review* 106 (2), 387–435.
- CBO (Congressional Budget Office). 2020. "The Distribution of Household Income, 2017." Tech. rep. CBO, Washington, DC.
- Eaton, Jonathan, and Harvey S. Rosen. 1980. "Labor Supply, Uncertainty, and Efficient Taxation." *Journal of Public Economics* 14 (3), 365–374.
- Feenberg, Daniel, and Elisabeth Coutts. 1993. "An Introduction to the TAXSIM Model." *Journal of Policy Analysis and Management* 12 (1), 189–194.
- Feldstein, Martin S. 1969. "The Effects of Taxation on Risk Taking." *Journal of Political Economy* 77 (5), 755–764.
- García-Miralles, Esteban, Nezih Guner, and Roberto Ramos. 2019. "The Spanish Personal Income Tax: Facts and Parametric Estimates." *SERIEs* 10 (3), 439–477.
- Goebel, Jan, Markus M. Grabka, Stefan Liebig, Martin Kroh, David Richter, Carsten Schröder, and Jürgen Schupp. 2019. "The German Socio-Economic Panel (SOEP)." *Jahrbücher für Nationalökonomie und Statistik* 239 (2), 345–360.
- Gourieroux, Christian, Alain Monfort, and Alain Trognon. 1984. "Pseudo Maximum Likelihood Methods: Applications to Poisson Models." *Econometrica* 52 (3), 701–720.
- Gouveia, Miguel, and Robert P. Strauss. 1994. "Effective Federal Individual Income Tax Functions: An Exploratory Empirical Analysis." *National Tax Journal* 47 (2), 317–339.
- Grabka, Markus M. 2020. "SOEP-Core v35 — Codebook for the \$PEQUIV File 1984–2018: CNEF Variables with Extended Income Information for the SOEP." Tech. rep., SOEP Survey Papers. Berlin.
- Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura. 2014. "Income Taxation of U.S. Households: Facts and Parametric Estimates." *Review of Economic Dynamics* 17 (4), 559–581.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2014. "Consumption and Labor Supply with Partial Insurance: An Analytical Framework." *American Economic Review* 104 (7), 2075–2126.

- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2017. "Optimal Tax Progressivity: An Analytical Framework." *Quarterly Journal of Economics* 132 (4), 1693–1754.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2020. "How Should Tax Progressivity Respond to Rising Income Inequality?" *Journal of the European Economic Association* 18 (6), 2715–2754.
- Heathcote, Jonathan, and Hitoshi Tsujiyama. 2021. "Optimal Income Taxation: Mirrlees Meets Ramsey." *Journal of Political Economy* 129 (11), 3141–3184.
- Holter, Hans A., Dirk Krueger, and Serhiy Stepanchuk. 2019. "How Do Tax Progressivity and Household Heterogeneity Affect Laffer Curves?" *Quantitative Economics* 10 (4), 1317–1356.
- Jakobsson, Ulf. 1976. "On the Measurement of the Degree of Progression." *Journal of Public Economics* 5 (1), 161–168.
- Kaas, Leo, Georgi Kocharkov, Edgar Preugschat, and Nawid Siassi. 2021. "Low Homeownership in Germany — A Quantitative Exploration." *Journal of the European Economic Association* 19 (1), 128–164.
- Kaplan, Greg. 2012. "Inequality and the Life Cycle." *Quantitative Economics* 3 (3), 471–525.
- Keane, Michael P. 2011. "Labor Supply and Taxes: A Survey." *Journal of Economic Literature* 49 (4), 961–1075.
- Kurnaz, Musab, and Terry A. Yip. 2022. "The Canadian Income Taxation: Statistical Analysis and Parametric Estimates." *Canadian Journal of Economics* 55 (1), 272–311.
- Li, Wenli, and Pierre-Daniel Sarte. 2004. "Progressive Taxation and Long-Run Growth." *American Economic Review* 94 (5), 1705–1716.
- Musgrave, Richard A., and Tun Thin. 1948. "Income Tax Progression, 1929–48." *Journal of Political Economy* 56 (6), 498–514.
- Park, Rolla E. 1966. "Estimation with Heteroscedastic Error Terms." *Econometrica* 34 (4), 888.
- Persson, Mats. 1983. "The Distribution of Abilities and the Progressive Income Tax." *Journal of Public Economics* 22 (1), 73–88.
- Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913–1998." *Quarterly Journal of Economics* 118 (1), 1–41.
- Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2017. "Distributional National Accounts: Methods and Estimates for the United States." *Quarterly Journal of Economics* 133 (2), 553–609.
- Rainey, Carlisle. 2017. "Transformation-Induced Bias: Unbiased Coefficients Do Not Imply Unbiased Quantities of Interest." *Political Analysis* 25 (3), 402–409.
- Schröder, Carsten, Johannes König, Alexandra Fedorets, Jan Goebel, Markus M. Grabka, Holger Lüthen, Maria Metzger, Felicitas Schikora, and Stefan Liebig. 2020. "The Economic Research Potentials of the German Socio-Economic Panel Study." *German Economic Review* 21 (3), 335–371.
- Silva, J. M. C. Santos, and Silvana Tenreiro. 2006. "The Log of Gravity." *Review of Economics and Statistics* 88 (4), 641–658.

- Varian, Hal R. 1980. "Redistributive Taxation as Social Insurance." *Journal of Public Economics* 14 (1), 49–68.
- Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge, MA.
- Wu, Chunzan. 2021. "More Unequal Income but Less Progressive Taxation." *Journal of Monetary Economics* 117, 949–968.