

Centre interuniversitaire de recherche
en économie quantitative

CIREQ

Cahier 10-2012

*On the Optimality of
Progressive Income Redistribution*

Ozan BAKIS and
Baris KAYMAK



Le **Centre interuniversitaire de recherche en économie quantitative (CIREQ)** regroupe des chercheurs dans les domaines de l'économétrie, la théorie de la décision, la macroéconomie et les marchés financiers, la microéconomie appliquée et l'économie expérimentale ainsi que l'économie de l'environnement et des ressources naturelles. Ils proviennent principalement des universités de Montréal, McGill et Concordia. Le CIREQ offre un milieu dynamique de recherche en économie quantitative grâce au grand nombre d'activités qu'il organise (séminaires, ateliers, colloques) et de collaborateurs qu'il reçoit chaque année.

*The **Center for Interuniversity Research in Quantitative Economics (CIREQ)** regroups researchers in the fields of econometrics, decision theory, macroeconomics and financial markets, applied microeconomics and experimental economics, and environmental and natural resources economics. They come mainly from the Université de Montréal, McGill University and Concordia University. CIREQ offers a dynamic environment of research in quantitative economics thanks to the large number of activities that it organizes (seminars, workshops, conferences) and to the visitors it receives every year.*

Cahier 10-2012

On the Optimality of Progressive Income Redistribution

Ozan BAKIS and Baris KAYMAK

Ce cahier a également été publié par le Département de sciences économiques de l'Université de Montréal sous le numéro (2012-09).

This working paper was also published by the Department of Economics of the University of Montreal under number (2012-09).

Dépôt légal - Bibliothèque nationale du Canada, 2012, ISSN 0821-4441

Dépôt légal - Bibliothèque et Archives nationales du Québec, 2012

ISBN-13 : 978-2-89382-632-5

On the Optimality of Progressive Income Redistribution*

Ozan Bakış

Barış Kaymak[†]

Galatasaray University and GIAM

University of Montreal and CIREQ

Markus Poschke[‡]

McGill University and CIREQ

Abstract

We compute the optimal non-linear tax policy for a dynastic economy with uninsurable risk, where generations are linked by dynastic wealth accumulation and correlated incomes. Unlike earlier studies, we find that the optimal long-run tax policy is moderately *regressive*. Regressive taxes lead to higher output and consumption, at the expense of larger after-tax income inequality. Nevertheless, equilibrium effects and the availability of self-insurance via bequests mitigate the impact of regressive taxes on consumption inequality, resulting in improved average welfare overall. We also consider the optimal once-and-for-all change in the tax system, taking into account the transition dynamics. Starting at the U.S. status quo, the optimal tax reform is slightly more *progressive* than the current system.

J.E.L. Codes: E20, E62, H21, H20

Keywords: Intergenerational Mobility, Optimal Taxation, Progressive Redistribution, Incomplete Markets

*We thank Ömer Açikgoz, Mark Bills, Rui Castro, Jonathan Heathcote, Remzi Kaygusuz, Dirk Krueger, Ananth Seshadri, Gianluca Violante, HakkıYazıcı and the seminar participants at the NBER Summer Institute, SED, New York/Philadelphia Quantitative Macro Workshop, University of Rochester, Sabanci University, Universitat Autònoma de Barcelona and Universitat de Barcelona for their comments.

[†]Department of Economics, Université de Montréal, C.P. 6128 succursale Centre-ville, Montréal, QC H3C 3J7, e-mail: baris.kaymak@umontreal.ca

[‡]Department of Economics, McGill University. e-mail: markus.poschke@mcgill.ca

1 Introduction

Most modern governments implement a redistributive fiscal policy, where incomes are taxed at an increasingly higher rate, while the provision of public services and transfers is skewed towards the poor. Such policies are thought to deliver a more equitable distribution of income and welfare, and, thereby, provide social insurance for future generations, who face uncertainty about what conditions they will be born into.

In market economies, such egalitarian policies can be costly as they disrupt the efficiency of resource allocation. Therefore, the added benefit of a publicly provided social safety net, that is over and above what is available to people through other sources, such as their family or the private sector, has to be carefully weighed against this cost. In this paper, we provide such an analysis of the optimal degree of income redistribution for a government that aims to maximize average welfare.

The possibility of a bad start is perhaps the most important economic risk in life. Studies show that 60-90% of the cross-sectional wage dispersion is explained by permanent differences among workers when they start their careers (Keane and Wolpin, 1997; Haider, 2001; Storesletten, Telmer, and Yaron, 2004), suggesting that pre-market factors, be they innate or acquired, are extremely important for subsequent economic success. Private provision of insurance against such risk naturally fails in this context. Since insurance, by definition, excludes pre-existing conditions, private provision would only be possible if a third party, such as parents, were able to sign their kids into obligations on their future earnings before they were born, or of legal age. Furthermore, unlike transitory shocks, permanent differences are the hardest to insure by individual means when markets are incomplete. These limitations may generate a case for publicly provided insurance through a redistributive tax system.

The optimal design of a redistributive tax system is, however, subject to constraints. We emphasize three. First, although a market for private insurance does not exist in our context, agents may have access to insurance through other means. Parental transfers,

in particular, provide a natural source of insurance against adverse economic outcomes. In order to prepare for risks faced by their offspring, parents accumulate precautionary funds. A redistributive tax policy would alleviate the need for parental insurance and crowd out accumulation of capital, leading to reduced investment.

Second, informational frictions may prevent the government from observing individual productivity. Consequently, it levies taxes on total income only, which leads to well-known incentive problems as higher taxes discourage workers from labor and thereby reduce output.

Third, the policymaker has to be cognizant of the implications of its tax policy on prices. Large-scale shifts in labor supply and savings alter the wage rate and the interest rate, which may have redistributive repercussions for income.

We address these constraints explicitly in a dynastic general equilibrium model with incomplete markets and endogenous labor supply, where generations are linked through a correlated income process. Families are not allowed to sign contracts contingent on their offspring's income. They can save, nonetheless, and transfer wealth to subsequent generations. They may not, however, pass their debt onto them. This is essentially an Aiyagari-Bewley-Hugget setting at dynastic frequency *à la* Barro (1974).

In this setting, we search for the optimal redistributive income tax scheme. Our approach to the problem is primarily quantitative and is in the tradition of Ramsey (1927).¹ The planner may not modify the financial structure of the economy. It cannot, for instance, introduce new assets or allow parents to accept obligations for their kids. It may, however, implement a transfer scheme, for example to transfer income to poor agents. Transfers and government expenditures are financed by taxes levied on labor and capital income. The set of tax policies is restricted to parametric forms albeit flexible ones. The tax schedule used here not only provides a good fit to the current U.S. system, but also

¹A parallel set of papers study the implications of information frictions in dynamic economics for allocations that are efficient under incentive-compatibility constraints (Mirrlees, 1971; Golosov, Kocherlakota, and Tsyvinski, 2003; Kocherlakota, 2005; Albanesi and Sleet, 2006; Farhi and Werning, 2012). The problem studied here is also one of constrained optimality, although the constraints are different.

allows for a variety of tax systems, such as progressive, flat, regressive and lump-sum.² We assume that the government can commit to a once-and-for-all change in the tax policy, and ask two questions: Which tax policy maximizes average welfare *at the steady state* of our model economy? Which tax policy maximizes average welfare *along a transition path*, starting from the current wealth and income distributions in the U.S.?

We find that the optimal tax policy for the long-run steady state is moderately regressive. When the government spends 25% of total output, the bottom fifth of the income distribution is taxed, on average, by 42%, and the top fifth pays 15% of their income, with a median tax rate of 29%. By contrast the current tax code in the U.S. calls for a tax rate of 16% tax for the lowest fifth, and 30% for the top fifth of the income distribution, with a median tax rate of 23%.

The intuition for this result is simple. A less progressive tax system fosters creation of wealth and income by raising the after-tax return to labor and savings, resulting in higher average consumption. The improvement in consumption levels is weighed against larger wealth and income inequality brought about by regressive taxation, an undesirable feature for a utilitarian government. The latter, however, is mitigated for two reasons. First, the larger supply of capital lowers the interest rate, while boosting the wage rate as labor complements capital in production. Consequently, the equilibrium price adjustment redistributes income away from the wealthy, who rely primarily on capital income, to consumption-poor agents who rely heavily on labor income, and counterbalances the increase in inequality generated by regressive taxation. Second, the availability of self-insurance through parental savings considerably limits the impact of income inequality on consumption inequality. These mechanisms are effective until moderate levels of regressivity, after which the marginal value of leisure outweighs that of additional income, which prevents further increases in hours worked. Output and average consumption stop rising, while inequality keeps growing, leaving no incentive for the government to

²Note that lump-sum taxes are not trivially optimal in our framework since the competitive equilibrium in Aiyagari (1994) is not constrained-efficient (Davila et al., forthcoming).

reduce progressivity any further.

When the short run dynamics are considered, a sudden switch to a regressive tax system from the current U.S. system is neither desirable nor politically feasible. Accumulation of the additional capital requires limited consumption of goods and leisure along the transition path, which limits the welfare gains from changing the tax policy. Furthermore, a sudden change in the tax system involves large and immediate transfers of income which generates substantial income inequality in the short-run. Due to discounting, these concerns outweigh the long-run benefits of regressive income taxation. A utilitarian government therefore prefers a tax system that is slightly more progressive than the current status quo in the U.S. once transition dynamics are taken into account.

The literature on optimal taxation is vast. The approach here is closest to Conesa and Krueger (2006) and Conesa, Kitao, and Krueger (2009), who calculate the optimal progressivity of income taxes for an OLG economy with incomplete markets and heterogeneous agents. Heathcote, Storesletten, and Violante (2010) take a similar approach to compute optimal progressivity in a Blanchard-Yaari-Bewley economy with partial insurance, and without capital. We differ crucially from these papers by allowing dynasties to self-insure via capital accumulation and bequests, and by introducing a correlation of income risk across generations. The results show that both components are important in gauging the value added by publicly provided social insurance, and for modeling the appropriate consumption response to tax policy. In particular, when insurance from non-public sources is available, a benevolent planner may prefer to improve social welfare by affecting incentives in the private insurance market and by harnessing general equilibrium effects rather than by directly providing insurance via income redistribution. As a result, the optimal tax schedule computed here is less progressive.

Erosa and Koreshkova (2007), Seshadri and Yuki (2004) and Benabou (2002) also look at taxation problems in dynastic settings, with emphasis on human capital investment and education. Benabou (2002) abstracts from dynastic capital accumulation and Se-

shadri and Yuki (2004) from labor supply. Both Erosa and Koreshkova (2007) and Shadri and Yuki (2004) analyze consequences of a flat tax reform, but do not calculate optimal non-linear taxation. Conesa and Krueger (1999) and Nishiyama and Smetters (2007) focus on payroll taxes and social security. Cutler and Gruber (1996), Rios-Rull and Attanasio (2000), Golosov and Tsyvinski (2007) and Krueger and Perri (2011) study how publicly provided insurance schemes can crowd-out insurance that is available through other sources. Hubbard, Skinner, and Zeldes (1995), in particular, emphasize the crowding out of precautionary savings by public tax policy.

2 A Dynastic Model with Redistributive Income Taxation

The model is a standard model of savings with uninsured idiosyncratic income risk (Aiyagari, 1994; Bewley, 1986; Hugget, 1993) applied at a dynastic frequency and extended to allow for non-linear fiscal policy and endogenous labor supply.

The economy consists of a continuum of heterogeneous consumers, a representative firm, and a government. We interpret each model period as a generation. There is a continuum of agents in a generation, each endowed with dynastic capital, k , and labor skill, z . With these endowments, they can generate an income of $y = zwh + rk$, where w is the market wage per skill unit, $h \in (0, 1)$ is hours worked and r is the interest rate net of depreciation.

Agents pay taxes on their income to finance an exogenous stream of government expenditure, which we assume proportional to aggregate output: $g_t = \gamma Y$. The disposable income of an agent net of taxes is given by $y^d(y)$, which depends only on the agent's total income. This function also determines the distribution of the tax burden.

Agents can allocate their resources between consumption and investment in dynastic capital, which can be used to transfer wealth to their offspring, but not to borrow from them. They derive utility from consumption, and they dislike work. They care about their welfare as well as their offspring's, which depends on the amount of wealth passed on by

their parent as well as their own skill endowment. The latter is determined stochastically by a first-order Markov process: $F(z'|z)$.

The problem of an agent is to choose labor hours, consumption and capital investment to maximize utility. The wage rate, the interest rate and the aggregate distribution of agents over wealth and productivity, denoted by Γ , are given. Let $\Gamma' = H(\Gamma)$ describe the evolution of the distribution over time. The Bellman equation for a consumer's problem is:

$$V(k, z; \Gamma) = \max_{c, k' \geq 0, h \in (0,1)} \left\{ \frac{c^{1-\sigma}}{1-\sigma} - \theta \frac{h^{1+\epsilon}}{1+\epsilon} + \beta \mathbb{E}[V(k', z'; \Gamma')|z] \right\} \quad (1)$$

subject to

$$\begin{aligned} c + k' &= y^d(y) + k \\ \Gamma' &= H(\Gamma). \end{aligned}$$

The production technology of a representative firm uses aggregate capital, K , and labor, N , as inputs, and takes the Cobb-Douglas form: $F(K, N) = K^\alpha N^{1-\alpha}$. Factor markets are competitive, and firms are profit maximizers.

A *competitive equilibrium* of the model economy consists of a value function, $V(k, z; \Gamma)$, factor supplies, $k'(k, z; \Gamma)$ and $h(k, z; \Gamma)$, a wage rate, $w(\Gamma)$, an interest rate $r(\Gamma)$, and an evolution function $H(\Gamma)$ such that:

- (i) Given $w(\Gamma)$, $r(\Gamma)$ and $H(\Gamma)$, $V(k, z; \Gamma)$ solves the worker's problem defined by (1) with the associated factor supplies $k'(k, z; \Gamma)$ and $h(k, z; \Gamma)$.

- (ii) Factor demands are given by the following inverse equations:

$$\begin{aligned} r(\Gamma) &= \alpha(K/N)^{\alpha-1} - \delta \\ w(\Gamma) &= (1-\alpha)(K/N)^{1-\alpha} \end{aligned}$$

(iii) Markets clear:

$$K = \int k'(k, z) d\Gamma(k, z) \quad \text{and} \quad N = \int zh(k, z) d\Gamma(k, z).$$

(iv) $H(\Gamma)$ is consistent with $F(z'|z)$ and the savings policy $k'(k, z; \Gamma)$.

(v) The government budget is balanced:

$$g = \int [y - y^d(y)] d\Gamma(k, z).$$

A *steady-state* of the economy is a competitive equilibrium where the distribution of agents is stationary, i.e. $\Gamma^{ss} = H(\Gamma^{ss})$.

2.1 A Redistributive Income Tax Policy

Taxes are modeled after the current U.S. income tax system, which can be approximated by a log-linear form for disposable income:

$$y^d = \lambda(zwh + rk)^{1-\tau}. \tag{2}$$

The power parameter $\tau \leq 1$ controls the degree of progressivity of the tax system, while λ adjusts to meet the government's budget requirement. When $\tau = 0$, the equation above reduces to the familiar proportional tax (or flat tax) system. When $\tau = 1$, all income is pooled, and redistributed equally among agents. For more moderate values, when $0 < \tau < 1$, the tax system is progressive.³

The disposable income function above also allows for negative taxes. Income transfers are, however, non-monotonic in income. When taxes are progressive, transfers are first increasing, and then decreasing in income. Examples of such transfers schemes include,

³The average tax rate is $1 - \lambda y^{-\tau}$, which is increasing in y if $\tau > 0$.

earned income tax credit, welfare-to-work programs etc. In Section 3, we show that this functional form provides a remarkable fit to the U.S. tax system.

A regressive tax system is achieved when τ is negative. In this case taxes are first increasing, then decreasing in income for high enough income levels, and may prescribe positive transfers for high income earners. Since the marginal tax rate, $1 - \lambda(1 - \tau)y^{-\tau}$, is monotonic in pre-tax income, (2) rules out tax policies that are progressive for some parts of the income distribution and regressive elsewhere.

In addition to the tax policies spanned by (2), a lump-sum tax system is considered as a benchmark. Note that a lump-sum system is not necessarily optimal since competitive equilibrium of the Aiyagari (1994) model is not constrained efficient (Davila et al., forthcoming).

2.2 Planner's Problem

The government is run by a benevolent planner who seeks to maximize average welfare in the economy. The planner chooses the progressivity of the tax policy subject to a balanced budget constraint and equilibrium responses by households to the tax policy. We consider two experiments. In the first one, the planner is concerned with the average welfare at the long-run steady-state of the economy. Formally, the problem is:

$$\max_{\lambda, \tau} \mathcal{W}^{ss} = \int V^{ss}(k, z; \Gamma^{ss}) d\Gamma^{ss}(k, z)$$

subject to

$$\begin{aligned} g &= \int [y - y^d(y; \lambda, \tau)] d\Gamma(k, z) \\ y &= wz h(k, z; \Gamma^{ss}) + rk'(k, z; \Gamma^{ss}). \end{aligned}$$

where V^{ss} is the value function, Γ^{ss} is the stationary distribution of agents over productivity and wealth, $h(\cdot)$ and $k'(\cdot)$ are the policy functions at the steady-state equilibrium

associated with the tax policy (λ, τ) . The dependence of these functions on the tax policy is suppressed for notational convenience.

In addition to the steady-state equilibrium, the planner may be concerned with the short-term consequences of a tax reform during the transition to the new steady state. Suppose that the government can credibly commit to a *once-and-for-all* change in the tax policy. In the second policy experiment, the planner seeks to maximize average utility by choosing the parameters of a tax reform. Specifically, the planner solves the following recursive problem:

$$\mathcal{W}(\Gamma) = \max_{\lambda, \tau} \int U(c(k, z; \Gamma), h(k, z; \Gamma)) d\Gamma(k, z) + \beta \mathcal{W}(\Gamma')$$

subject to

$$\begin{aligned} g &= \int [y - y^d(y; \lambda, \tau)] d\Gamma(k, z) \\ y &= wz h(k, z; \Gamma) + rk'(k, z; \Gamma) \\ \Gamma' &= H(\Gamma), \end{aligned}$$

where $H(\Gamma)$ satisfies the consistency condition of the competitive equilibrium and the starting distribution Γ_0 is given.

3 Empirical Analysis and Calibration

The model is calibrated to the U.S. economy with emphasis on intergenerational income risk.⁴ The model period is 25 years. The discount parameter β is calibrated to generate an annual interest rate of 4.3%. The capital share of income, α , is set to 0.36, and the depreciation rate to 8% per annum. The rate of relative risk aversion is set to 2.0.

This leaves three sets of parameters: the fiscal policy, (γ, λ, τ) , the preference parame-

⁴The findings are qualitatively similar when life-cycle income fluctuations are considered. Results from a calibration that incorporates life-cycle income risk to the model are presented in the appendix.

ters for labor, θ and ϵ , and the stochastic intergenerational income process, $F(z'|z)$. These parameters are identified as follows.

3.1 How Progressive is the U.S. Tax System?

The progressivity of the current tax system is estimated using household-level data from March supplements to the Current Population Survey for 1979 to 2009. Federal and state income taxes, as well as the payroll tax per household are obtained from the NBER tax simulator (Feenberg and Coutts, 1993). Our measure of pre-tax income is gross earnings, as reported by the household, plus the payroll tax. Disposable income is defined as reported earnings less federal and state income taxes. The estimated log-linear regression is:

$$\log y^d(y) = 1.34 + 0.83 \log y + X\hat{\Gamma} \quad R^2 = 0.94,$$

where X includes indicators for survey year. The correlation coefficient indicates that the log-linear specification fits the U.S. tax system remarkably well. Figure 1 further confirms this visually by plotting average disposable income by quantiles of pre-tax income (circles) over the regression line (solid). Two points are worth noting. First, the slope of the regression line is less than one, showing the progressivity of the U.S. tax system. The implied value of τ is 0.17 (0.0026).⁵ This also seems to be a good approximation of the progressivity of the tax system defined over life-time income, as modeled here.⁶ Second, the bottom five percent of the gross-income distribution are paying negative or zero taxes.

[Figure 1 about here.]

⁵Corporate taxes are not available in our dataset. To test the relevance of this for our estimate, we estimated the same specification for 2004 based on the information in Table 2 of Piketty and Saez (2007), who impute corporate taxes in their calculations using federal tax returns. We estimate the progressivity to be 0.164, virtually the same as our estimate above.

⁶To compare the cross-sectional estimate of τ to one that would apply to life-time income, we simulated earnings processes based on data from the PSID, and re-estimated τ . The resulting estimate was close to the one used here with the difference being less than 0.01.

We estimate the size of the government expenditures as 25% of output based on total income taxes (state and federal) and payroll taxes relative to GDP. Given $\tau = 0.17$ and $\gamma = 0.25$, the value of λ is determined at the equilibrium by the government's budget constraint.

3.2 Intergenerational Wage Mobility

The stochastic intergenerational income process is estimated directly using data on multiple generations from the PSID (1968 - 2009). First, a fixed effects regression is estimated for hourly wages for respondents of ages 24 to 60, controlling for indicators for age and survey year. Then fathers are paired with sons to estimate the intergenerational wage transition matrix.⁷

[Table 1 about here.]

The results are shown in Table 1. The transition matrix displays a significant degree of persistence. The implied average intergenerational correlation of wages is 0.35 (0.004), which is close to the values reported in the literature. The last row shows the average wage rate in each quartile of the life-time wage distribution in 1999 dollars.⁸

3.3 Leisure and Labor Supply

The preference parameter for labor disutility, θ , is calibrated to average hours worked over life for a generation. The curvature of the utility with respect to hours worked, ϵ , governs two crucial moments in the model economy: the intergenerational elasticity of labor substitution, and the cross sectional dispersion of hours worked within a generation. Since little is known about the former, we calibrate this parameter to the coefficient of variation of average lifetime labor hours. In the next section, we test our calibration using

⁷For fathers with multiple sons, we replicate the wage observations for the father.

⁸All income values were adjusted by CPI.

the model’s implications for the intergenerational correlation of hours, and by calculating a pseudo-elasticity across generations.

To calculate the cross-sectional distribution of life-time hours within a cohort, we calculate a similar fixed effects regression, controlling for age and survey year. To capture variations in labor force participation (e.g. due to retirement, schooling or cyclical variations), all observations on workers of ages 15 to 75 are included in the regression. Those years where the respondent reported working zero hours are also included. An average person in the sample works 1,908 hours a year, which makes up 44% of their available time.⁹ The coefficient of variation is 0.29.

3.4 Calibration Results

Table 2 summarizes the calibrated values for the parameters. The implied values for the utility parameters are: $\theta = 0.305$, $\epsilon = 0.825$, and $\beta = 0.965$. Next, we evaluate the calibration results by comparing the predictions of the model for labor supply elasticities, intergenerational correlations of household wealth and hours.

[Table 2 about here.]

Labor Supply Elasticity

The elasticity of labor supply is a crucial parameter of interest for gauging the distortionary effects of taxation on hours. In the model, the labor supply elasticity depends on the tax policy and the prices in the economy. These are kept fixed at their calibrated values to compare the model’s predictions for individual labor supply schedules with the estimates found in the literature.

Since the model is dynastic, any change in the wage rate is, by construction, permanent over the lifetime of the generation. The relevant measure of elasticity to gauge the changes in the labor supply is the Marshallian elasticity. In the benchmark calibration,

⁹The total available time is calculated as 4,368 hours (= 12 hours $\times$ 7 days $\times$ 52 weeks).

the (Marshallian) wage elasticity varies from 0.07 to 0.58, with an average of 0.42, and a standard deviation of 0.10. As expected, it is decreasing in wealth, and increasing in productivity. The uncompensated (pre-tax) income elasticity is -0.72 on average with a standard deviation of 0.10. It is decreasing in wealth, and non-monotonic in productivity. These values are well within the range of estimates reported in Blundell and MaCurdy (1999).

The utility function features a constant Frisch elasticity of $(1 - \tau_l)/(\epsilon + \tau_l)$, which equals 0.84 given the benchmark calibration. It is hard to compare this to the estimates in the literature since the intertemporal substitution of labor in a dynastic context is not a well-defined concept. Nonetheless, the micro level estimates for yearly models are around 0.25 for individuals, while a value between 2 and 3 is required to match employment differences across time and countries at the macro level (See Prescott (2004); Cho and Cooley (1994) and Blundell and MaCurdy (1999) among others.)

Intergenerational Correlations

The model is calibrated to the intergenerational mobility of wages. Given the optimal policies for labor supply and savings, the model has implications for the persistence of hours and wealth from one generation to the next. Table 3 reports the transition matrix for wealth computed in the model, and compares it with the one estimated by Charles and Hurst (2003) using data from the PSID. While the model captures the main features of wealth transitions, it predicts somewhat higher persistence, especially for the top quintile. The strength of persistence is more pronounced in the transition of hours. The elasticity of offspring's working hours with respect to parent's working hours is 0.85 in the model, whereas it is around 0.20 in the data.

[Table 3 about here.]

The model is expected to produce more persistent wealth and hours distributions, because the agents in the model differ only with respect to their productivity. Their pref-

erences towards consumption and labor are identical, whereas the data may contain other dimensions of heterogeneity. Disutility of labor, for instance, may vary from one member of a dynasty to the next. Second, the wealth transitions in Charles and Hurst (2003) are based on wealth before the parent generation is deceased, i.e. it excludes the final bequest, and should be considered as a lower bound on persistence of wealth. Finally, the data may contain measurement error, which leads to a seemingly more mobile transition matrix.

4 How Progressive Should the Long-Run Tax Policy Be?

To determine the optimal tax policy, we run two tax policy experiments. In the first experiment, the steady-state equilibrium is computed for different tax policies with varying degrees of progressivity. In the second experiment, we compute the expected welfare along a transition path to a steady-state equilibrium, in response to a once-and-for-all change in the progressivity of the tax code. In this section, we focus on the first experiment, and discuss the optimal long-run tax policy and its projected impact on the economy. Then we conduct counterfactual experiments to isolate the role of different modeling assumptions on our results.

The findings suggest that the optimal tax code in the long-run is moderately regressive. The optimal value for τ is -0.38. The tax rates implied by this value are shown in Table 4. The optimal long-run tax policy calls for an average tax rate of 42% for the lowest decile of the income distribution compared to 15% for the top decile. The median tax rate is 29%. By contrast, the benchmark economy, calibrated to the U.S. tax policy, has an average tax rate of 16% for the bottom decile relative to 30% for the top decile. The median tax rate in the benchmark economy is 23%.

[Table 4 about here.]

Although average tax rates are monotonically declining in income at the optimal steady-

state, taxes are not. Tax payments are increasing in income for the first two thirds of the income distribution, and begins to decline for the top tertile. The share of taxes paid by the lowest income decile is 6% compared to 9% for the top decile.

How could a regressive tax system, which subjects low income groups to higher tax rates, be optimal for an egalitarian government? To see this, note that a utilitarian policymaker is concerned with two things when comparing tax policies: the total amount of available goods (consumption and leisure), and how these goods are distributed among agents. A less progressive tax policy raises the average level of consumption at the cost of higher *after-tax* income inequality. This does not translate to equally severe consumption inequality since agents self-insure via dynastic capital accumulation. As a consequence, the optimal tax schedule may well be regressive if the increase in inequality is sufficiently small compared to the increase in average consumption.

Table 5 summarizes the impact of progressivity on the steady-state of the model economy. The first column reports the values for the benchmark economy calibrated to the U.S. tax policy ($\tau = 0.17$). Columns further right display less progressive tax systems, and the last column shows the optimal tax code.

[Table 5 about here.]

A decline in the progressivity of the tax policy promotes generation of income by increasing the after-tax return to labor and capital. This raises savings in the economy. Less progressive income redistribution also raises the risk faced by future (unborn) generations, and gives parents an incentive to accumulate additional precautionary savings. For high-income groups, there is an additional income effect generated by lower taxes, which further encourages accumulation of capital. For low-income groups, the income effect works against the substitution effects, but is not strong enough. Overall, supply of capital increases, which puts a downward pressure on the interest rates.

The larger capital stock has two implications for labor. First, it raises the demand for labor, and increases the wage rate, despite the downward pressure created by the increase

in the labor supply. Second, larger wealth has a negative income effect on labor supply, limiting the increase in labor input, and pushing the wage rate further up. With a larger stock of capital and increased labor input, output increases. The optimal tax system leads to a 44% increase in output, which translates to a 34% increase in consumption. The rise in welfare due to higher average consumption is slightly mitigated by the decrease in average leisure from 0.56 to 0.48. A more important mitigating factor is the rise in inequality, which we turn to next.

4.1 Tax Progressivity and Inequality

Overall, an average person in an economy with less progressive taxes has larger wealth, higher income, substantially more consumption, and slightly less leisure. To compare this improvement in the utility of an average person with the potential changes in distributive inequality, Table 6 shows the Gini coefficients for crucial variables in the model. The economy with regressive taxes features larger wealth inequality along with a considerable increase in the inequality of after-tax income disposable for consumption. The Gini coefficient for wealth inequality increases from 0.51 to 0.59, and from 0.16 to 0.25 for disposable income. The latter is roughly equal to the increase in income inequality in the U.S. during the second half of 20th century.

[Table 6 about here.]

The impact of rising income and wealth inequality on consumption, however, is limited. The Gini coefficient for consumption inequality rises from 0.13 to 0.18, about half the rise in disposable income inequality. This is due, in large part, to the availability of self-insurance through dynastic capital. This result is also consistent with Krueger and Perri (2006), who find that the rise in consumption inequality has been muted relative to income inequality after 1980.

The Gini coefficient for leisure inequality increases from 0.16 to 0.20. This is even less than the increase in the inequality of goods consumption. Overall, the rise in welfare

inequality remains small relative to the gains in average consumption.

The change in equilibrium prices also helps alleviate the effects of declining progressivity on pre-tax income inequality. The decline in the interest rate mitigates the effect of rising wealth inequality on income inequality, while the higher wage rate increases the weight on labor income, which is more equally distributed under regressive taxes. These help explain the relatively stable pre-tax income inequality in Table 6.

4.2 Tax Progressivity and Steady-State Welfare

The improvement in average steady-state welfare when the economy switches to the optimal tax code can be measured in consumption units for comparison. To calculate a consumption equivalence, we ask the following hypothetical question: by what factor would one need to increase the consumption of each and every person in the benchmark economy to reach the same average welfare as the optimal economy, keeping their labor supply constant? The answer is 9.1%. Such an improvement in welfare is quite large considering that the welfare cost of business cycles are estimated at 1% or less.¹⁰ This calculation ignores the changes in welfare during the transition to the new steady-state, which is studied in Section 5.

[Figure 2 about here.]

To see how the distribution of welfare across agents changes, we first compare the value functions for a given wealth and productivity level, without taking into account the shift in the wealth distribution. Figure 2 plots welfare by wealth for the lowest and the highest productivity groups (out of 4 in total). The solid lines correspond to the benchmark economy, and the dashed lines represent the economy operating under the optimal tax code. The optimal economy features lower welfare for the wealthy, especially for

¹⁰For a risk aversion of two (as here), Krebs (2007) reports 0.98%, which is much larger than the estimate reported in Lucas (1987).

those with little labor income. This is primarily due to the lower interest rate in the optimal economy. Workers with low wealth, on the other hand, are dependent on labor income, which is higher in the new economy due to higher wage rates. This leads to higher welfare for the highly productive, who have higher disposable incomes in the new tax system, and mitigates the fall in welfare for workers with low productivity and, hence, income, who receive less transfers.

A utilitarian policymaker also considers the shift in the wealth and income distributions when comparing these two economies. In particular, the optimal economy features a higher wealth level on average, which leads to an upward movement along the dashed welfare functions in Figure 2.

4.3 Labor Supply, Self-Insurance and Partial Equilibrium: Implications for Tax Policy

We emphasized three crucial constraints on the policymaker's choice of redistributive tax policy: the crowding out of labor supply, availability of self-insurance via parental wealth and adjustment of prices in equilibrium. To highlight the relative roles of these constraints for the optimality of progressive redistribution, three counterfactual calculations are presented in this section.

We conduct two experiments to gauge the implications of dynastic wealth for optimal taxes. First, we recompute the optimal tax code assuming that savings behavior is fixed at the benchmark economy. This prevents agents from optimally adjusting their precautionary savings policy to the tax system. It therefore shuts down the response of private insurance to public insurance. Prices and labor supply are allowed to respond optimally, and the budget is balanced at all times. Shutting down the savings response to tax policies prevents accumulation of new capital in response to less progressive taxes. Consequently, less progressive tax policies lead to consumption inequality without any improvement in aggregate output and consumption. As a result, the optimal tax policy in this case is moderately progressive with τ of 0.27.

The experiment above shuts down the savings response, but is not comparable to models without capital. To emphasize the role of capital, we simulate a second counterfactual economy, where capital is held entirely by the government, and supplied competitively to firms. The government keeps the total stock of capital constant at the U.S. benchmark level. The return on capital is deducted from total tax obligations of workers. Workers have no wealth, hence, there are no savings decisions to be made. They choose their labor supply every period. Essentially, this is a static model. The optimal tax policy in this hypothetical economy features considerable progressivity with a τ of 0.46. This result is not too surprising. Relative to the previous scenario, workers are now stripped of their ability to retain any wealth to insure themselves against income fluctuations. Any shock to income is a shock to consumption, raising the need for insurance through a tax system. This results in a more progressive optimal tax code. The only factor that prevents full redistribution from being optimal in this economy is the endogenous labor supply.

These two experiments show that the interaction between public and private insurance is key for the results above. When the private insurance (or precautionary savings) response to changes in the tax system is shut down, regressive taxes don't induce capital accumulation and are not optimal anymore.

The optimal tax system also implies strong price effects. To gauge their importance, we compute optimal taxes for a partial equilibrium economy, where the wage rate and the interest rate remain at their benchmark levels. Savings and labor supply still respond optimally, and the government runs a balanced budget. With fixed prices, the changes in savings and labor supply in response to a decline in progressivity do not translate into higher wages and lower interest rates. This shuts down the redistributive role of equilibrium price adjustment, making regressive taxes less attractive. The optimal taxes are moderately progressive in this case with a τ of 0.23.

If, in addition, the labor supply response is shut down in the partial equilibrium economy, the optimal progressivity increases from 0.23 to 0.31. Since redistribution is not

allowed to crowd out labor supply, optimal policy becomes more progressive. At the same time, the increase in progressivity is modest, partly because lifetime labor supply is not very elastic.

These experiments reveal that two features are key for the optimality of a tax system: the effect of taxes on private insurance and capital accumulation, and the resulting changes in prices in general equilibrium. Therefore, the presence of capital in the analysis is crucial.

The findings here could also be viewed in the light of the findings in Davila et al. (forthcoming), who show that in Aiyagari (1994) type models, agents' individual savings decision in the laissez-faire equilibrium imply that the economy does not reach its constrained efficient optimum. When the income of the poor consists mainly of labor income, as here, the utilitarian planner prefers to subsidize savings to promote capital accumulation, which raises the wage rate and, hence, income for the low consumption groups. The planner in Davila et al. (forthcoming) implements this policy by state-dependent tax and transfer schemes on capital income. When state-specific policies are not feasible, a regressive income tax system can stand-in for more complicated mechanisms.

Also note that the results are not driven by joint taxation of capital and labor. In a version of the model with a linear tax on capital income and a non-linear tax on labor income of the form shown in equation (2), the optimal progressivity of labor income taxation is -0.26 when the tax rate on capital is 25%, and -0.23 when capital is not taxed at all.

5 Optimal Redistribution along a Transition Path

The optimal tax code described in the previous section encourages capital accumulation and accordingly leads to substantially higher wages than the benchmark economy. Getting there is costly, however: Increased capital accumulation requires initially reducing consumption and/or leisure. Therefore, the transition to the steady state following a switch to a regressive tax system is costly and matters for welfare. Comparing steady

states abstracts from this cost. Depending on its size, implementing the tax code that is optimal in the long run may not be optimal once the transition is taken into account. Overall welfare including the transition may instead be maximized by a different tax code. Therefore, we next ask the following two questions: What is the welfare effect of implementing the tax code that is optimal at the steady state of the economy? And which level of progressivity of the tax code is optimal, taking into account the transition from the current U.S. benchmark?

5.1 Transition to the Optimal Steady State

In analyzing this issue, we assume that the economy initially is in the benchmark steady state that reproduces the U.S. status quo. In this situation, the government surprisingly implements the new tax code and commits to it. As the economy converges to the new steady state induced by the changed tax system, the interest rate, the wage rate and λ all change. Recall that government expenditure is a constant fraction of output, and the parameter λ of the tax code adjusts to balance the government's budget every period. The algorithm used to compute the transition path is described in the appendix.

[Figure 3 about here.]

Results show that the transition to the optimal long run policy is very costly. Values of key endogenous variables along the transition path are shown in Figure 3. The economy moves into the neighborhood of the new steady state in about 4 periods. Over this time, the capital stock is more than doubled and consumption is increased by a third. Early in the transition, however, increased capital accumulation implies much lower leisure. Furthermore, a sudden change in the tax policy brings about a substantial increase in the after-tax income inequality, and, thereby, consumption inequality, for the first generation. Unlike the rise in average consumption, which is realized in the future, the increase in consumption inequality is immediate. Since each period lasts 25 years, these early periods

have very high weight. (Subsequent generations carry much lower weight, e.g. the weight of the fifth generation is only about 1%.) Therefore, while the only generation that actually loses in terms of V is the first one, its loss is so large and its weight so heavy that the transition becomes undesirable. As a consequence, the cost of the transition wipes out the welfare gains achieved in the steady state with the regressive tax policy, and it is not optimal to implement that policy.

Could a regressive tax reform raise enough political support? This depends on the share of winners from such a policy. Suppose that there were a referendum at the benchmark steady state on the optimally regressive tax policy. Only those who are economically active during that period can vote in the referendum. The value functions along the transition path indicate that only 30% of the population would vote in favor of a regressive tax policy. This support consists of mostly (currently) high income earners and the wealthy. The percent of support by productivity (in an increasing order) is 1%, 3%, 13% and 82%. A similar picture emerges across wealth groups. Less than 1% of those with low wealth (lowest third) support the policy, compared to 30% of agents in the middle wealth group and 57% of the wealthiest third.

Implementing the transition to the optimal steady state is not optimal because losses suffered by early generations outweigh the benefits received by later generations. This raises the question whether support for the reform could be garnered if side payments from winners to losers were permitted.¹¹ For the transition to the optimal steady state, consumption of period-1 winners could be reduced by 3.3% of status quo output without dropping their welfare below the status quo level. However, raising the welfare of period-1 losers to the status quo level would require increasing their consumption by 12.4% of status quo output. Side payments therefore are not sufficient to make this reform acceptable. This is similar for other large reforms. The most regressive reform that could be acceptable with side payments is a transition to τ of 0.05, which would require payouts

¹¹A similar scheme is applied by Nishiyama and Smetters (2007) in their analysis of social security privatization.

of just below 1.3% of status quo output to compensate losers, while winners would still be better off after facing a consumption reduction of 1.3% of status quo output.

5.2 Optimal Tax Reform along the Transition

This raises the question which tax reform is optimal, starting in the U.S. status quo. We find that the optimal reform of progressivity consists in slightly increasing progressivity, from 0.17 to 0.26.

Comparing steady states, higher progressivity results in a 17% decrease in the capital stock, with most of the decrease taking place in the first three periods. It also implies slightly higher leisure, resulting in a 3% decline in the aggregate labor input. These two changes induce a fall of about 7% in aggregate consumption. Applying the same criterion as in Section 4.2, agents in the status quo would require about 4% larger consumption in every state to be willing to live instead in the steady state implied by the more progressive policy.

[Figure 4 about here.]

Why then is this an optimal reform? The answer lies in the transition. Apart from the difference in direction, the transition path, which is shown in Figure 4, is qualitatively similar to the transition to the optimal steady state. A progressive tax reform leads to an immediate redistribution of income, leading to lower inequality of income and consumption for the current generation, a desirable outcome for a utilitarian planner. Increased progressivity also discourages labor supply, leading to higher average leisure for the current generation. Due to discounting, the combination of these two desirable features outweighs the loss of average consumption by future generations.

Overall, 48% of the agents are better off under a progressive tax reform. Those with higher productivity or higher wealth fare worse. The level of political support is decreasing in wealth (81%, 51% and 13% by tertiles of wealth) and in productivity (95%, 80%, 46% and 0% for four productivity states).

6 Discussion

The results highlight the different roles of public policy as an insurance device. When alternative sources of insurance are available, the planner may find it optimal to harness the effect of taxes on demand for private insurance and, thereby, on the equilibrium prices to improve social welfare. A regressive tax system achieves this by raising the wage rate and lowering the interest rate, giving the policymaker the option of eliminating the distortionary effects of a progressive system while keeping inequality under control. When the long-run outcomes of public policy are considered, this alternative is preferable to direct provision of social insurance via progressive taxation. The latter proves more costly as it also crowds out labor supply and savings.

By contrast, direct redistribution through a progressive tax reform is preferable when short-run dynamics are considered. Immediate reduction of inequality outweighs a reduction in long-run welfare. This is due to two reasons. First, arriving at a steady-state with a higher capital stock requires longer labor hours along the way. Second, since capital is inelastically supplied in the short-run, the policymaker may find it optimal to raise the tax rate for high income groups, who rely more heavily on capital income. This is akin to the result in Chamley (1986), where the government finds it optimal to initially confiscate the entire private capital stock before lowering the tax rate on capital to zero.

If the tax policy were announced ahead of time, the planner would not be able to redistribute the existing resources as effectively. The capital supply would decline in response to the announcement before the policy takes effect. As a result, the planner would prefer a tax reform that is less progressive. Moreover, when announced sufficiently early, future redistribution of resources is less costly for the current generation since the future state of a dynasty does not depend strongly on its current state (unless shocks are fully persistent). This helps increase the political support for a tax reform.

Weighting of generations is also crucial for the optimality along the transition. If, for instance, the planner assigned (almost) equal weights to all generations, the welfare at the

steady-state would outweigh the short-run considerations. An alternative is to design intergenerational tax and transfer schemes to bring part of the increase in long-run average welfare to the current generation to compensate for longer hours. An implementation of such transfers may be possible if the planner can issue current debt against payments by future generations. How the possibility of debt affects the optimal progressivity of income tax policy remains as a promising venue for future research.

Appendix

A Life-Cycle Income Risk and Optimal Progressivity

The focus has so far been on the roles of intergenerational income risk and dynastic wealth accumulation for optimal taxation, whereas income fluctuations over the life-cycle were assumed away. This choice was motivated by the fact that life-cycle income risk constitutes a much smaller fraction of total income inequality. In this section, we check the sensitivity of the findings to this assumption.

To do this, we set the time period to 5 years and introduce a probability of death μ of 0.20 per period.¹² If an agent dies, his successor draws a new productivity level given the transition matrix in Table 1, and takes over the dynastic capital. If the agent survives, which happens with probability $1 - \mu$, he draws a new productivity level from a different transition matrix that governs income fluctuations over a worker's life.¹³ We assume perfect altruism. The rest of the model remains the same.

We estimate the life-cycle income risk using the PSID by calculating the transition probabilities for 5-year intervals. The first panel in Table 7 shows the estimated probabilities. The levels of productivity are defined by the distribution of average lifetime wages as in the previous section. Compared to the intergenerational transitions, there is much

¹²The period length is 25 years in the benchmark model and 5 years in the lifecycle model, which gives a relative duration of 5 periods. A death probability of 20% yields an average life of 5 periods.

¹³A somewhat similar approach was employed in Castaneda, Diaz-Gimenez, and Rios-Rull (2003).

more persistence over life. Conditional on being born in the first quartile, the chances that a worker moves up the quartile ladder in the next 5 years is 25%. Conditional on survival, it is 56% after 25 years, our period length in the benchmark specification. Across generations, this probability is 66%.

[Table 7 about here.]

To see the impact of adding life-cycle income risk, the second panel in Table 7 shows the combined transition matrix over 25 years assuming that the probability of death is 0.20. Overall, introducing the life-cycle income process reduces the persistence of income shocks, especially for higher income categories.

We recalibrate the model to the same targets as in Table 2, which changes two of our parameters: the disutility of labor, θ , which is now 0.34, and the parameter of labor supply elasticity, ϵ , which becomes 0.7. Remaining parameters do not change. Under this calibration, we find that the optimal tax policy is slightly regressive with $\tau = -0.09$, which is close to a flat tax system. The associated welfare gain from switching to the optimal tax code is equivalent to a 2.69% rise in consumption.

When the model is extended to allow for life-cycle income risk, two competing effects on optimal progressivity arise. On the one hand, agents find it easier to self-insure against income shocks that are now less persistent. Consequently, there is less motivation for insurance through progressive taxes. On the other hand, with lower persistence, the precautionary savings motive weakens. This mutes the savings response to regressive income taxes, and limits the response of the capital stock and, thereby, aggregate consumption. In our extended calibration, average consumption increases by 15%, and the coefficient of variation for consumption increases from 0.22 to 0.26. Both are much smaller compared to the benchmark economy, where average consumption goes up by 30%, and the inequality increases to 0.32.

References

- Aiyagari, Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics* 109 (3):659–684.
- Albanesi, S. and C. Sleet. 2006. "Dynamic Optimal Taxation with Dynamic Optimal Taxation with Private Information." *Review of Economic Studies* 73 (1):1–30.
- Barro, R. 1974. "Are Government Bonds Net Wealth?" *Journal of Political Economy* 82 (6):1095 – 1117.
- 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.
- Bewley, T. 1986. *Contributions to Mathematical Economics: In honor of Gerard Debreu*, chap. Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers. New York: North Holland, 79 – 102.
- Blundell, R. and T. E. MaCurdy. 1999. *Labor Supply: A Review of Alternative Approaches*, chap. 27. Handbook of Labor Economics. North Holland, Amsterdam, 1559 – 1694.
- Caroll, Christopher D. 2006. "The method of endogenous gridpoints for solving dynamic stochastic optimization problems." *Economics Letters* 91 (3):312–320.
- Castaneda, A., J. Diaz-Gimenez, and J. V. Rios-Rull. 2003. "Accounting for the U.S. Earnings and Wealth Inequality." *Journal of Political Economy* 111 (4):818–857.
- Chamley, C. 1986. "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives." *Econometrica* 54 (May):607–22.
- Charles, Kerwin and Erik Hurst. 2003. "The Correlation of Wealth across Generations." *Journal of Political Economy* 111 (6):1155–1182.

- Cho, J. and T. F. Cooley. 1994. "Employment and Hours Over the Business Cycle." *Journal of Economic Dynamics and Control* 18 (2):411 – 432.
- Conesa, J. C. and D. Krueger. 1999. "Social Security Reform with Heterogeneous Agents." *Review of Economic Dynamics* 2 (4):757 – 795.
- Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger. 2009. "Taxing Capital? Not a Bad Idea After All!" *The American Economic Review* 99 (1):25–48.
- Conesa, Juan Carlos and Dirk Krueger. 2006. "On the optimal progressivity of the income tax code." *Journal of Monetary Economics* 53 (7):1425–1450.
- Cutler, D.M. and J. Gruber. 1996. "Does public insurance crowd out private insurance?" *The Quarterly Journal of Economics* 111 (2):391.
- Davila, J., J. H. Hong, P. Krussel, and V. Rios-Rull. forthcoming. "Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks." *Econometrica* .
- Erosa, Andres and Tatyana Koreshkova. 2007. "Progressive taxation in a dynastic model of human capital." *Journal of Monetary Economics* 54 (3):667–685.
- Farhi, E. and I. Werning. 2012. "Capital Taxation: Quantitative Explorations of the Inverse Euler Equation." *Journal of Political Economy* 120 (3):398–445.
- Feenberg, Daniel Richard and Elizabeth Coutts. 1993. "An Introduction to the TAXSIM Model." *Journal of Policy Analysis and Management* 12 (1):189–194.
- Golosov, M., N. Kocherlakota, and A. Tsyvinski. 2003. "Optimal Indirect and Capital Taxation." *Review of Economic Studies* 70 (3):569–87.
- Golosov, M and A Tsyvinski. 2007. "Optimal Taxation with Endogenous Insurance Markets." *Quarterly Journal of Economics* 122 (2):487–534.

- Haider, S. J. 2001. "Earnings Instability and Earnings Inequality of Males in the United States: 1967–1991." *Journal of Labor Economics* 19 (4):799 – 836.
- Heathcote, Jonathan, Kjetil Storesletten, and Gianluca Violante. 2010. "Redistributive Taxation in a Partial-Insurance Economy." Mimeo.
- Hubbard, R Glenn, Jonathan Skinner, and Stephen P Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy* 103 (2):360–399.
- Hugget, M. 1993. "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies." *Journal of Economic Dynamics and Control* 17 (5-6):953–69.
- Keane, Michael P. and Todd I. Wolpin. 1997. "The Career Decisions of Young Men." *Journal of Political Economy* 105 (3):473–522.
- Kocherlakota, N. 2005. "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation." *Econometrica* 73 (5):1587 – 1621.
- Krebs, T. 2007. "Job Displacement Risk and the Cost of Business Cycles." *American Economic Review* 97 (3):664 – 686.
- Krueger, D. and F. Perri. 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory." *Review of Economic Studies* 73:163 – 193.
- Krueger, Dirk and Fabrizio Perri. 2011. "Public versus private risk sharing." *Journal of Economic Theory* 146:920–956.
- Lucas, Robert E. Jr. 1987. *Models of Business Cycles*. New York: Blackwell.
- Mirrlees, James A. 1971. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38 (2):175–208.
- Nishiyama, S. and K. Smetters. 2007. "Does Social Security Privatization Produce Efficiency Gains?" *Quarterly Journal of Economics* 122 (1):1677 – 1719.

- Piketty, Thomas and Emmanuel Saez. 2007. "How Progressive is the U.S. Federal Tax System? A Historical and International Perspective." *Journal of Economic Perspectives* 21 (1):3–24.
- Prescott, E. J. 2004. "Why Do Americans Work So Much More than Europeans?" *Federal Reserve Bank of Minneapolis Quarterly Review* 28 (1):2–13.
- Ramsey, F. 1927. "A Contribution to the Theory of Taxation." *Economic Journal* 37:47 – 61.
- Rios-Rull, JV and O Attanasio. 2000. "Consumption smoothing in island economies: Can public insurance reduce welfare?" *European Economic Review* 44:1225–1258.
- Seshadri, Ananth and Kazuhiro Yuki. 2004. "Equity and efficiency effects of redistributive policies." *Journal of Monetary Economics* 51:1415–1447.
- Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron. 2004. "Consumption and Risk Sharing over the Life Cycle." *Journal of Monetary Economics* 51 (3):609–633.

Table 1: Intergenerational Wage Mobility in the U.S.

Quartile	1	2	3	4
1	0.34	0.33	0.17	0.16
2	0.18	0.31	0.34	0.17
3	0.12	0.29	0.31	0.28
4	0.07	0.16	0.24	0.53
wage rate (\$/hr)	9.7	14.9	20.2	33.0

Note.— Transition matrix between average lifetime (ages 24 to 60) wages of father-son pairs, controlling for age and year effects. Data comes from the PSID (1968 - 2009).

Table 2: Calibration of the Model to the U.S. Economy

<i>Parameter</i>	<i>Value</i>	<i>Target Moment</i>	
T	= 25	1/3 of life-time	
σ	= 2.00	relative risk aversion	2.00
$\beta^{1/T}$	= 0.965	annual interest rate	4.3%
θ	= 0.305	average annual labor hours	0.44
ϵ	= 0.825	coef. of variation of hours	0.29
α	= 0.36	capital share of income	0.36
δ	= 0.88	annual depreciation rate	0.08
γ	= 0.25	authors' estimates of G/Y	
τ	= 0.17	authors' estimates	

Table 3: Intergenerational Transition of Wealth

	1	2	3	4	5		1	2	3	4	5
1	0.57	0.26	0.17	0.00	0.00	1	0.36	0.26	0.16	0.15	0.11
2	0.43	0.33	0.11	0.13	0.00	2	0.29	0.24	0.21	0.13	0.16
3	0.07	0.28	0.34	0.31	0.00	3	0.16	0.24	0.25	0.20	0.14
4	0.00	0.04	0.38	0.26	0.33	4	0.12	0.15	0.24	0.26	0.24
5	0.00	0.00	0.02	0.31	0.67	5	0.07	0.12	0.15	0.26	0.36

(a) Model

(b) PSID

Note.— Transition probabilities by quintiles of the wealth distribution. The PSID estimates are taken from Table 2 of Charles and Hurst (2003).

Table 4: Average Tax Rates by Income

Tax Policy	Income Percentiles						
	<0.10	0.10–0.25	0.25–0.50	0.50–0.75	0.75–0.90	0.90–0.99	>0.99
Average Tax Rates (%)							
Benchmark (U.S.)	16	18	22	26	29	29	30
Flat Tax	25	25	25	25	25	25	25
Lump-sum	42	35	30	24	18	18	17
Optimal	42	37	32	26	16	16	15
Total Tax Share (%)							
Benchmark (U.S.)	1.5	7	21	28	25	16	2.1
Flat Tax	3.0	11	22	22	27	16	1.7
Lump-sum	7.3	18	24	24	15	11	1.1
Optimal	6.1	18	26	25	15	9.3	0.9

Table 5: Optimal Tax System: Steady-State Comparison

	Benchmark (U.S.)	Flat Tax	Lump Sum	Optimal
Progressivity (τ)	0.17	0.00	–	–0.38
Output	14.1	16.1	19.7	20.3
Interest Rate (%)	4.27	3.68	2.88	2.74
Wage Rate	0.21	0.22	0.25	0.26
Wealth/Capital	1.87	2.47	3.72	3.98
Hours Worked	0.44	0.46	0.51	0.52
Labor Input	44.0	46.2	50.4	50.8
Pre-tax Income	12.5	14.0	16.5	16.8
Consumption	9.4	10.5	12.3	12.6
Welfare	-0.255	-0.245	-0.239	-0.238

Table 6: Progressivity of Tax Policy and Inequality

	Benchmark (U.S.)	Flat Tax	Lump Sum	Optimal
Pre-tax Income	0.19	0.19	0.16	0.18
Disposable Income	0.16	0.19	0.22	0.25
Wealth	0.51	0.54	0.55	0.59
Consumption	0.13	0.15	0.16	0.18
Leisure	0.16	0.18	0.19	0.20
Welfare	0.14	0.16	0.19	0.20

Notes.— Table shows the Gini coefficient of inequality.

Table 7: Wage mobility and Income Risk

	1	2	3	4
1	0.75	0.17	0.05	0.04
2	0.27	0.43	0.21	0.08
3	0.10	0.22	0.41	0.25
4	0.04	0.05	0.14	0.80

(a) Life-Cycle

	1	2	3	4
1	0.33	0.23	0.18	0.26
2	0.29	0.22	0.19	0.30
3	0.25	0.21	0.20	0.34
4	0.21	0.19	0.20	0.40

(b) Life-Cycle and Intergenerational

Note.— Transition matrices for wages over the life-cycle (a) and in combination with intergenerational transmissions (b).




Figure 1: The progressivity of the U.S. tax system – Disposable household income as a function of pre-tax income. Circles denote the average income net of taxes by quantiles of gross income. The solid line is the fitted regression line. Data combines March supplements to CPS (1979 - 2009) with the NBER tax simulator.




Figure 2: Welfare by Wealth and Productivity




Figure 3: Transition to Regressive Taxation




Figure 4: Optimal Tax Policy along the Transition