Federal Reserve Bank of Minneapolis
Research Department

December 2019

Quantifying Efficient Tax Reform*

Job Boerma
University of Minnesota

Ellen R. McGrattan
University of Minnesota
and Federal Reserve Bank of Minneapolis

ABSTRACT

This paper quantifies welfare gains from Pareto reforms in an overlapping generations framework with policies constrained due to private information about shocks to household labor productivity. We use administrative panel data for the Netherlands to first estimate key parameters under status quo policies for households in different education groups. We then solve for Pareto optimal reforms and decompose the source of welfare changes into gains from level effects and gains from improved insurance. For our baseline parameterization, we find large welfare gains, on the order of 20 percent of lifetime consumption. Optimal consumption allocations are higher and smoother than allocations under current policy, while leisure allocations are lower and more volatile. We show that the welfare decomposition is quantitatively sensitive to estimates of wage profiles and processes governing shocks to labor productivity.

Keywords: income tax reform

JEL classification: E13, H21

*Boerma acknowledges support from the Heller-Hurwicz Economic Institute, and McGrattan acknowledges support from the NSF. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Introduction

This paper quantifies welfare gains from Pareto reforms in an overlapping generations framework with policies constrained due to private information about shocks to household labor productivity. We use administrative panel data for the Netherlands to first estimate key parameters under status quo policies for households in different education groups. We then solve for Pareto optimal reforms and decompose the source of welfare changes into gains from level effects and gains from improved insurance.

To model the Netherlands, we use a small open economy framework with overlapping generations and households that are heterogeneous in age, education, and productivity. Fiscal policy in this economy is summarized by tax schedules on incomes and assets and a tax rate on consumption. We compute values under current policy and use them—along with estimates for preferences, technologies, and wage processes—as inputs to our reform problem. In the reform problem, we compute the maximum consumption equivalent gain, which is the same for all households.

For our baseline parameterization, we find large welfare gains, on the order of 20 percent of lifetime consumption. Optimal consumption allocations are higher and smoother than allocations under current policy, while leisure allocations are lower and more volatile. This is to be expected given the planner is providing more insurance. To investigate this further, we decompose the total gain into contributions for level effects and contributions for improved insurance—for consumption and leisure. Increasing mean consumption is by far the largest source of gain, although some education groups with high variability in wages also have significant gains in lowering consumption dispersion.

We also show that the welfare decomposition is quantitatively sensitive to estimates of wage profiles and processes governing shocks to labor productivity. We explore two variations of the baseline model. First, we turn off growth in wages over the life cycle. In this case, the gains from smoothing consumption are close to zero, even for households with significant variation in their labor productivity shocks. Second, we lower variances of shocks for all households. Here again, we find a significant effect on estimates of the gains for lowering dispersion in allocations.

This paper is related to the literature on optimal income taxation. We extend Farhi and Werning (2013) and Golosov, Troshkin, Tsyvinski (2016) to compute Pareto reforms using a baseline matched to the Netherlands and allow for more general productivity shocks. Like Hosseini and Shourideh (2019), we compute the set of Pareto improving policy reforms, but we allow for stochastic productivity shocks. We find that allowing for stochastic shocks is quantitatively important for our welfare decomposition.

2. Theory

In this section, we describe the positive economy that is our baseline for estimating key parameters of preferences, technologies, and current fiscal policies of an actual economy. We then describe the associated planning problem used to quantify Pareto reforms of the original OLG economy.

2.1. Positive Economy

In this section, we describe the model economy that will be matched up to administrative data for the Netherlands. The environment is relatively standard with the exception of country-specific fiscal policies. There are a large number of households facing uninsurable productivity risks and perfectly competitive firms with constant-returns-to-scale technologies.

Households differ by age j, assets a, and productivity ϵ. They solve the following dynamic program:

$$v_j\left(a,\epsilon;\Omega\right)=\max_{c,n,a'}\left\{U\left(c,\ell\right)+\beta E[v_{j+1}\left(a',\epsilon';\Omega\right)|\epsilon]\right\}$$

subject to the budget constraint

$$a'=\left(1+r\right)a-T_a\left(ra\right)+w\epsilon n-T_n\left(j,w\epsilon n\right)-\left(1+\tau_c\right)c$$

and a lower bound on asset holdings: $a'\geq 0$. The aggregate state vector contains prices and policies:

$$\Omega=\{r,w,G,B,T_a,T_n,\tau_c\},$$

where r is the interest rate, w is the wage rate, G is government consumption, B is government debt, $T_a(\cdot)$ is the tax schedule for financial assets, $T_n(\cdot)$ is the tax schedule for labor income less transfers, and τ_c is the tax rate on consumption.

We assume the economy is small and open with interest rates r set in international markets. Firm technologies are constant-returns-to-scale functions in capital K and labor N with output given by:

$$Y=F\left(K,N\right).$$

Thus, knowing r, we also know the aggregate capital-labor ratio K/N and the wage rate w from the firm's optimality conditions. For computations below, we assume that prices and policies are fixed over our sample.[1] In a competitive equilibrium, the resource constraint

[1] This assumption can be easily relaxed without adding much computational burden as shown by Nishiyama and Smetters (2014).

must also hold in all periods:

$$C_t + K_{t+1} - (1-\delta) K_t + G_t + B_{t+1} - RB_t = Y_t.$$

The key outputs obtained from computing equilibria for the positive economy are the values under current policy, namely

$$\vartheta\left(\epsilon^{j-1}\right) = E\left[v_j\left(a, \epsilon; \Omega\right)|\epsilon_-\right]$$

or, in the case of future generations $\vartheta(\epsilon_0) = E[v_1(a, \epsilon; \Omega)|\epsilon_0]$. We want to Pareto improve on these value, ideally by shifting the allocations to the efficient frontier. As an example, consider the two-person case drawn in Figure 1. The allocation for the Netherlands (point NL) is the result of calculating the equations above. In the next section, we compute a reform problem that puts the two households on the efficient frontier, at a point with their consumption levels higher by the same percentage (say, Δ).

2.2. Reform Problem

In this section, we describe the planning problem that we solve to compute Pareto reforms given the initial valuations from the positive economy that will later be matched to administrative data from the Netherlands.

As in the positive economy, we assume that the interest rate r is given as we are working with an open economy. Given the aggregate production function $F(K, N)$, we can infer the capital-labor ratio K/N and the wage rate w. We also assume that the planner must finance government spending G and takes the initial assets B_0 as given.

Given these initial values for all households, the planning problem is to choose a feasible allocation that maximizes excess initial resources so that remaining lifetime values exceed their initial values for all households. Formally, the planning problem is:

$$\max\ F\left(K_0, N_0\right) + RB_0 - C_0 - K_1 + (1-\delta) K_0 - G_0 - B_1$$

subject to the laws of motion for capital and the resource constraints for all periods as in the positive economy, along with incentive constraints that ensure truthful reporting of households private productivity, and a condition such that lifetime value exceeds the given initial value. In the appendix, we prove that an allocation is Pareto efficient if and only if it solves the planning problem given the initial values.

It turns out that the Lagrange function for the planning problem is separable in the allocation of each household and, therefore, we can separately characterize the solution to the planning problem for each household. The planner problem for a household is to choose a household allocation to maximize excess resources subject to the household's incentive constraints. To make this tractable, we assume that only local downward incentive constraints bind at the solution. Assuming that only the local downward incentive constraints bind is a finite type analog for the first-order approach typically adopted in dynamic Mirrlees problems with a continuum of productivity types. (See, for example, Farhi and Werning (2013), Golosov, Troshkin, and Tsyvinski (2016), and Stantcheva (2017).) The relaxed component planner problem is formulated by replacing the set of constraints that ensure global incentive compatibility in the component planning problem with the set of constraints that ensure the allocation satisfies all local downward incentive constraints. We write the relaxed component planner problem recursively and then characterize its solution.

This problem can be formalized mathematically as follows. The planner chooses sequences of consumption $c_j(\epsilon)$, labor $n_j(\epsilon)$, promised values $V_j(\epsilon)$ for telling the truth about the productivity type, and threat values $\tilde{V}_j(\epsilon)$ for reporting a productivity type of ϵ while being one level more skilled, which we denote by ϵ^{+}. The recursive planning problem is given by:

$$\Pi_j\left(V_-, \tilde{V}_-, \epsilon_-\right) \equiv \max \sum_{\epsilon_i \in \mathcal{E}} \pi_j(\epsilon_i|\epsilon_-) \left(w_t \epsilon_i n_j(\epsilon_i) - c_j(\epsilon_i) + \Pi_{j+1}\left(V_j(\epsilon_i), \tilde{V}_j(\epsilon_{i+1}), \epsilon_i\right) / R\right)$$

subject to:

$$\begin{aligned} & U\left(c_j(\epsilon_i), \ell_j(\epsilon_i)\right) + \beta V_j(\epsilon_i) \\ & \qquad \geq U\left(c_j(\epsilon_{i-1}), \ell_j^{+}(\epsilon_{i-1})\right) + \beta \tilde{V}_j(\epsilon_i), i = 2, \ldots, N \end{aligned} \tag{2.1}$$

$$V_- = \sum_{\epsilon_i \in \mathcal{E}} \pi_j(\epsilon_i|\epsilon_-) \left(U\left(c_j(\epsilon_i), \ell_j(\epsilon_i)\right) + \beta V_j(\epsilon_i)\right) \tag{2.2}$$

$$\tilde{V}_- = \sum_{\epsilon_i \in \mathcal{E}} \pi_j\left(\epsilon_i|\epsilon_-^{+}\right) \left(U\left(c_j(\epsilon_i), \ell_j(\epsilon_i)\right) + \beta V_j(\epsilon_i)\right), \tag{2.3}$$

where $\pi_j(\epsilon_i|\epsilon_-^{+})$ is the conditional probability over current states ϵ_i for households that were one level more productive in the previous period ϵ_-^{+}. The first set of constraints in (2.1) ensures that utility is higher under truth-telling, with the leisure arguments given by:

$$\ell_j(\epsilon_i) = 1 - n_j(\epsilon_i)$$

$$\ell_j^{+}(\epsilon_{i-1}) = 1 - n_j(\epsilon_{i-1})\, \epsilon_{i-1}/\epsilon_i.$$

When calculating the welfare of efficient reform, we replace V_- in the problem above with $\vartheta(\epsilon_0) + \vartheta_\Delta$, where $\vartheta(\epsilon_0)$ is the initial value for future generations—that is, $E[v_1(0, \epsilon; \Omega)|\epsilon_0]$ in the positive economy—and the ϑ_Δ is the value of giving Δ more consumption to households.

3. Data and Estimation

In this section, we discuss the administrative data from the Netherlands and estimation methods used to parameterize the model. We start with aggregated data from their national accounts, flow of funds, population censuses, and tax authorities. We then discuss the micro data on earnings, hours, and education.

3.1. Aggregate Data

The main data source for the aggregate data is the Dutch Bureau of Statistics. These data are publicly available.

3.1.1. National accounts

The primary data source for national income and product accounts is the *nationale rekeningen*.[2] Table 1 splits national income by factor of production. Labor income includes compensation of employees and 70% of proprietors' income. All other income is categorized as capital income, which we adjust in three ways. First, we subtract product-specific taxes as measured in the government's income and expenditure accounts. We make this correction because we are interested in production at producer prices rather than at consumer prices. Second, we impute capital services for consumer durables—which we treat as investment—and government capital. The imputed services are assessed to be 4 percent of the current-cost net stock of consumer durables and government fixed assets. Government fixed assets as well as consumer durables are recorded as non-financial balanses. Finally, we impute depreciation of consumer durables. Since our data do not include the equivalent of the United States flow of funds, we assume the ratio of consumer durable depreciation

[2] The main table is *BBP vanuit de inkomensvorming* and is found under: macro-economie; nationale rekeningen; BBP, finale bestedingen en productie; BBP.

to consumer durable goods to be identical to the United States.[3] This implies consumer durable depreciation of 5 percent.

On the product side, revisions must also be made with regard to sales taxes, capital services and consumer durables depreciation. The sales taxes are assumed to primarily fall on personal consumption expenditures. We assume pro rata shares when assessing how much of the taxes are on durables, non-durables and services. We include nondurables and services with consumption and durable goods with tangible investment. Therefore, we subtract sales taxes from both product categories. Imputed capital services only affect our consumption measure, which combines personal and government consumption from the national accounts. The consumption of consumer durables depreciates the outstanding stock of durables, which motivates us to classify consumer durables depreciation as consumption.

Fixed assets and other capital stocks used in our analysis are shown in Table 2 with averages for 2000-2010. As in the case of national accounts, we divide all estimates by adjusted GDP. We add the stock of consumer durables. The data are separated for businesses, households, and the government. We also include the value of land, which is much higher than estimates reported by McGrattan and Prescott (2017) for the United States. In fact, the data show that the value of residential land exceeds the value of stuctures by roughly 12 percent, likely due to strict government regulation of land use. Since the oil and gas sector is so significant for the Netherlands, we include reserves. Related to fixed assets are the valuations in flow of funds data which we report in Table 3. Here, we report estimates for household net worth and government debt relative to GDP averaged over the sample 2000-2010.

[3] See Table 1 in McGrattan and Prescott (2017).

Finally, in Table 4, we report aggregates on population and hours, which we use to parameterize preferences and to check aggregated micro data. Averaging data between 2000 and 2010, we estimate that the Dutch population worked 12,243 million hours, implying average annual hours of 1,135 for every individual between ages 16 and 64.

The data from the national accounts and population census are used to parameterize the discount factor, the capital share in F, the depreciation rate, the weight on leisure in preferences, the length of working life, and the length of retirement.

3.1.2. Fiscal Policy

In Figure 2, we plot the income tax schedule for the Netherlands during our sample period. The figure shows three marginal tax rates, namely, 34, 42 and 52 percent for working age households, with cutoff levels of 20,000 and 59,000 euro. Marginal tax rates are reduced for retirees with incomes below 35,000 euro. Specifically, the marginal tax rate is 17 percent for incomes below 20,000 euro and 24 percent for incomes below 35,000 euro. In Figure 3, we show the tax schedule for financial assets. Below 46,000 euro, the tax rate is 0. Above this level, the rate of taxation is 1.2 percent. Finally, we assume an effective tax rate on consumption of 13.4 percent, which is the weighted average VAT for a basket of goods in the Netherlands. These schedules and rates are used to parameterize T_a, T_n, and τ_c.

3.2. Micro Data

We use linked administrative records between 2006 and 2014 from Statistics Netherlands for the information on education, earnings, and hours—series that we need to estimate productivity processes $\{\epsilon\}$ and wage profiles $\{\zeta\}$ over the life cycle for different education groups.

3.2.1. Merged datasets

We start with a representative subsample of all Dutch households selected by Statistics Netherlands. The sample consists of roughly 95 thousand households per year, which is 1.3 percent of the population of households, covering a total of over 275 thousand individuals. For all analyses, we weight households with the provided sample weights. We consider all households with heads of household above age 25. Income is measured by employer-provided earnings records. We construct an individual's annual taxable labor earnings, which includes the employer's health insurance contribution, by adding all earnings reports within a given calendar year. To construct an hourly wage rate, we merge the earnings dataset with a dataset on employer-reported hours worked, dividing taxable labor earnings by hours of work. Because the model features a single decision maker for each household, we define the household wage rate for married and cohabitating households as the average individual wage rate weighted by the hours worked of each partner. For single households, the individual wage rate is the household wage rate. Household non-market time is given by average individual non-market time which is discretionary time minus individual hours worked. We set an individual's discretionary time equal to 16 hours a day for 365 days.

We merge the datasets for earnings and hours with another that provides education levels for our sample. We need this information because we assume that there is ex-ante heterogeneity in productivity and wage profiles based on the highest educational degree earned. We classify every degree as a low, a medium, or a high level of education. The low education level is a high school degree or a practical degree, the medium level is a degree from a university of applied sciences, and the high level is a university degree. We group households into six education bins, which are unordered pairs of the degree of each partner. Singles are grouped with couples in which both partners have obtained the same level of education.[4]

We should note here that there are significant advantages to the merged data available in the Netherlands relative to what is available in most other countries. For example, in the case of the United States, we only have administrative data for earnings whereas in the Netherlands we have earnings and hours linked and available for all members of the household. We also have detailed data on education, which is not available in the United States.

3.2.2. Estimated wage processes

We estimate the parameters that govern the residual wage process using the minimum distance estimator (Chamberlain (1984)). We first regress logarithmic wages on as follows:

$$\log W_{ijt} = A_t + X_{ijt} + \omega_{ijt},$$

where the household index is i, age is j, and the period is t. The right-hand-side variables are time effects A_t and household observables X_{ijt}. One of the observables is the age of the head of household, the coefficient of which is our estimate of the lifecycle profile ζ_j.

The second step is to estimate components of the residual wage after pooling across cohorts. More specifically, we use the method of simulated moments approach to estimate

[4] In our sensitivity analysis, we also explore conditioning on head of household, which is more common in the literature.

parameters ρ, σ_u^2, σ_η^2, $\sigma_{\epsilon_0}^2$ for the standard permanent-transitory process:

$$\omega_{ij} = \epsilon_{ij} + \eta_{ij}$$

$$\epsilon_{ij} = \rho\epsilon_{ij-1} + u_{ij}$$

with the persistent component of the residual wages given by ϵ_{ij} and the transitory component given by η_{ij}. The error processes and initial conditions are assumed to be distributed normally, that is $\eta_{ij} \sim \mathcal{N}(0, \sigma_\eta^2)$, $u_{ij} \sim \mathcal{N}(0, \sigma_u^2)$, and $\epsilon_{i0} \sim \mathcal{N}(0, \sigma_{\epsilon_0}^2)$.

The moments we use to identify the parameters are the variances and first-order autocovariances. These moments can be written in closed form as follows:

$$\begin{aligned} \operatorname{var}(\omega_{ij}) &= \rho^{2j}\sigma_{\epsilon_0}^2 + \frac{1-\rho^{2j}}{1-\rho^2}\sigma_u^2 + \sigma_\eta^2 \\ \operatorname{cov}(\omega_{ij}, \omega_{ij-k}) &= \rho^k \frac{1-\rho^{2(j-k)}}{1-\rho^2}\sigma_u^2 + \rho^{2j-k}\sigma_{\epsilon_0}^2. \end{aligned}$$

These expressions are functions of (j, k) and the four parameters.

The estimation of the wage process uses the minimum distance estimator introduced by Chamberlain (1984), which minimizes a weighted squared sum of the differences between each moment in the model and its data counterpart. Let $m(\Lambda)$ be the vector of theoretical covariances and Λ be the parameter vector. The data counterpart is given by $\hat{m}$. In this case, the estimator solves:

$$\min_{\Lambda} \; (\hat{m} - m(\Lambda))' \, \mathcal{W} \, (\hat{m} - m(\Lambda)),$$

where $\mathcal{W}$ is a weighting matrix. For our baseline parameterization, we use the identity matrix for $\mathcal{W}$. To compute confidence intervals, we bootstrap using 1,000 replications. Given the closed form expressions for the theoretical moments, the objective function is efficiently evaluated.

We use the estimated parameters ρ and σ_u^2 to parameterize the residual wage process in the model.[5] The results of our estimation procedure are reported in Table 5. We find the parameters are precisely estimated with estimates for $\hat{\rho}$ in the range of 0.95 to 0.97 across education groups. If we construct estimates of variation for the residual wages, that is, $\hat{\sigma}_u^2/(1-\hat{\rho}^2)$, we find that households with a Low and High member and those with two High members are close to twice as variable as the others.

In Figure 4, we report the life-cycle wage profiles (ζ_j) for the 6 education groups. The right side of the figure shows the population for each group. For example, the Low-Low group is the largest with 43 percent of the working population. We have normalized these estimates by dividing each profile using the average wage for the entire population. Not surprising, we find a steep rise between ages 25 and 45 for all groups, with the lowest higher by roughly 40 percent and the highest by roughly 200 percent.

3.3. Computation

When we compute equilibria for our positive economy and our reform problem, we approximate labor productivity shocks by a Markov chain with 20 types. For both problems we assume that baseline preferences are logarithmic, that is,

$$U(c,\ell) = \gamma \log c + (1-\gamma)\log \ell. \tag{3.1}$$

Both problems are parallelizable and thus we can solve them quickly on most modern computer clusters.

[5] We assume η is a shock that households can insure against, and we use the ergodic distribution based on ρ and σ_u^2 to parameterize the initial distribution of productivities. One could also use $\sigma_{\epsilon_0}^2$.

4. Results

In this section, we report our main findings for the baseline model and several alternative parameterizations, namely, cases with double the number of types, no wage growth, and lower variances in the shock processes.

4.1. Baseline

The main deliverables for our baseline model are labor wedges and welfare gains. We compute labor wedges for each education group. These wedges represent distortions used by the planner to incentivize individuals and to provide insurance across time and across types. We then report consumption-equivalent welfare gains and their decomposition into gains from increasing the level of consumption, gains from reducing dispersion in consumption, gains from increasing the level of leisure, and gains from reducing the dispersion in leisure.

The labor wedges are defined as follows:

$$\tau_n\left(\epsilon^j\right) = 1 - \frac{1}{w}\frac{U_\ell\left(c\left(\epsilon^j\right), \ell\left(\epsilon^j\right)\right)}{U_c\left(c\left(\epsilon^j\right), \ell\left(\epsilon^j\right)\right)} \tag{4.1}$$

and computed for each education group. Equation (4.1) tells us that in the optimal allocation there is a wedge between the wage rate w and the marginal rate of substitution between consumption and leisure. We report these wedges for the baseline model in Figure 5. The highest wedge is not that of the High-High group, but rather the Low-High group. The reason is that the Low-High group has the most variable wage process. The greater the dispersion in productivities, the greater are gains from redistribution and higher is $\tau_n(\epsilon^j)$. In fact, if we were to take averages, we would find a positive correlation between the total variance $\hat{\sigma}_u^2/(1-\hat{\rho}^2)$ and the wedge across education groups. This information is useful for the reform of current policy.

In Table 6, we report the welfare gains and its decomposition for our baseline parameterization. We find a total gain of 20 percent for an efficient reform in which all individuals are made better off by the same percentage. Building on Floden (2001) and Benabou (2002), we decompose this total gain into the gain from increasing consumption, the gain from smoothing consumption, the gain from increasing leisure, and the gain from smoothing leisure. That is, we take the total consumption-equivalent gain Δ and compute:

$$\log(1+\Delta) = \log\left(\left(1+\Delta_c^L\right)\left(1+\Delta_c^D\right)\right) + (1-\gamma)\log\left(\left(1+\Delta_\ell^L\right)\left(1+\Delta_\ell^D\right)\right)/\gamma.$$

Let $\hat{x}$ be the allocation in the planner problem and x be the allocation in the positive economy. Then we define the gain due to a level increase in $x = c$ or $x = \ell$ as

$$1+\Delta_x^L = \frac{\sum \pi\left(\epsilon^j\right)\hat{x}\left(\epsilon^j\right)}{\sum \pi\left(\epsilon^j\right)x\left(\epsilon^j\right)}.$$

We define the gain due to a reduction in dispersion in $x = c$ or $x = \ell$ as:

$$1+\Delta_x^D = \sum \beta^j \pi\left(\epsilon^j\right)\log\left(\frac{\hat{x}\left(\epsilon^j\right)}{\sum \pi\left(\epsilon^j\right)\hat{x}\left(\epsilon^j\right)}\right) - \sum \beta^j \pi\left(\epsilon^j\right)\log\left(\frac{x\left(\epsilon^j\right)}{\sum \pi\left(\epsilon^j\right)x\left(\epsilon^j\right)}\right)$$

The results of the decomposition are shown in Table 6. First, note that the summing across rows yields the total gain of 20 percent (and may be off because of rounding). Second, note that there are large gains for increasing and smoothing consumption, but the optimal plan calls for lower and more dispersed leisure than in the positive economy. The gains from increasing consumption are the most significant. For the Low-High and High-High groups, there are also significant gains from reducing dispersion since the variances of wages are largest for these groups. The planner lowers leisure the most for the most productive. However, in terms of leisure dispersion, the most noteworthy group is High-High that face significantly more leisure dispersion under the optimal plan relative to the positive economy.

If we consider these results in light of more simple models—say, static models with and without insurance—we find that our results are in line with the simpler models. For example, consider the case in which there is no insurance and households maximize (3.1) subject to a budget constraint that consumption is less than or equal to after-tax labor earnings. The optimal plan in that case calls for variation in consumption but constant leisure. If instead there was full insurance, a planning problem would call for constant consumption and variation in leisure. The positive economy is closer to the no-insurance case and the reform problem is closer to the full-insurance case.

In Figure 6, we show the allocations of (log) consumption and leisure along with their variances for those in the Low-Low group—which accounts for 43 percent of the population. In the upper left panel of the figure, we have plotted consumption for ages 25 to 64. We see from the figure that the planner can completely smooth mean consumption, which is not possible under current policy. In the upper right panel of the figure, we have plotted the variance of consumption. Dispersion is lowered in early years, but is higher than the positive economy later in life. In the lower panels, we plot the results for leisure. As predicted, leisure is lower in the reform problem than the positive economy for most years. while the variance is higher.

4.2. Sensitivity

We recompute total consumption-equivalent gains and their decomposition in three alternate specifications. First, we double the number of types from 20 to 40. Second, we set the wage profiles in Figure 4 to 1 for all types and all ages. Third, we lower the variance σ_u^2 of the labor productivity shocks for all education groups. In each case, we compare results for the levels and dispersion of allocations across ages for the Low-Low group to the baseline model.

The first set of results are shown in Table 7. In this case, we find no change in overall welfare between the baseline model with 20 types and the alternative with 40 types. The main difference between models is in the attribution of gains for the high variance groups (Low-High and High-High). For these groups, there are larger gains to raising consumption and lower leisure in the 40-type model relative to the baseline model. For the Low-Low group, which accounts for most of the population, the differences are not large. This is evident in Figure 7 where we compare the allocations—both levels and variances—for the two models. We see almost no change in levels and small changes in the variances of the allocations.

If we set wage profiles equal to 1 for all ages and groups, we find a slightly lower overall gain of 18 percent when compared to the 20 percent gain in the baseline. In Table 8, we see that the main difference between these models is the attribution of gains for smoothing consumption. If the profiles are flat, there are no gains to smoothing consumption: they are already smooth. In the case of the Low-Low group, we see in Figure 8 that the positive economy consumption and leisure allocations are already quite flat.

Finally, we compare the baseline model to one with the labor productivity variance σ_u^2 lowered by two-thirds. The overall gain in the latter case is 18 percent, slightly lower than the baseline and, not surprising, the gains for reducing consumption and leisure dispersion are smaller than in the baseline. These results are reported in Table 9. In Figure 9, we again compare allocations to the baseline model for the Low-Low group. Here, we find no difference in the levels, but lower variances for consumption and leisure in the alternative model when compared to the baseline.

In all experiments, we find large gains to the Pareto reforms, but the sources of the gains depend importantly on the estimated processes for labor productivities.

5. Conclusion

In this paper, we computed efficiency gains of Pareto reforms in an environment with policies constrained due to private information about shocks to household labor productivity. Using administrative data for the Netherlands, we found the gains to be large but also found that quantifying the sources of these gains depends importantly on having precise estimates for household wage profiles and shock processes for labor productivity.

Appendix

In the main text we present a planning problem and discuss how we characterize its solution using relaxed and recursive household planning problems. In this appendix we describe the intermediate steps. Our discussion closely follows Boerma (2019).

A household identity is a birth year t along with a productivity history ϵ^{j-1}. We denote a household by $i \equiv (t, \epsilon^{j-1})$. The set of households is partitioned into households alive in the initial period and households born in future periods, $\mathcal{I} \equiv \left\{ \{(j-1, \epsilon^{j-1})\}_{j=1}^{J}, \{(t, \epsilon_0)\}_{t=1}^{\infty} \right\}$.

An allocation for household i is a sequence of functions that specify consumption and labor supply at age $j+v$ given the household's productivity history ϵ^{j+v}, $x(i) \equiv \{x_{t+v}(\epsilon^{j+v})\}_{v=0}^{J-j} = \left\{ \left(c_{t+v}(\epsilon^{j+v}), n_{t+v}(\epsilon^{j+v}) \right) \right\}_{v=0}^{J-j}$. An allocation x specifies an allocation for every household i and aggregate quantities:

$$x \equiv \left\{ \{x\,(i)\}_{\mathcal{I}}\,, \{(C_t, N_t, B_{t+1}, K_{t+1})\}_{t=1}^{\infty} \right\}.$$

An allocation is resource feasible if and only if the allocation satisfies the resource constraint and the law of motion for capital in all periods.

Households know their history ϵ^j at age j, and the only source of information about this history are reports by the household itself.[6] By the revelation principle we restrict the reporting space to be the type space without loss. We use $\sigma_j(\epsilon^j)$ to denote the report that the household gives about their age j shock when they experience ϵ^j. A reporting strategy, specifying a report for every history, is denoted $\sigma \equiv \{\sigma_j(\epsilon^j)\}_{\epsilon^j, j}$. A reporting strategy generates a report history $\sigma^j(\epsilon^j) = (\sigma_1(\epsilon^1), \ldots, \sigma_j(\epsilon^j))$. Let Σ be the set of

[6] To simplify the exposition we describe incentive compatibility for a household born in the future and we suppress the identity. The corresponding definitions for households that are alive in the initial period naturally follow.

reporting strategies. The truthful reporting strategy is such that $\sigma^j(\epsilon^j) = (\epsilon_1, \ldots, \epsilon_j)$ for all j and all $\epsilon^j \in \mathcal{E}^j$.

Given a reporting strategy, the corresponding household allocation is given by $x^\sigma \equiv \{x_j(\sigma^j(\epsilon^j))\}_{\mathcal{E}^j,j} = \left\{c_j(\sigma^j(\epsilon^j)), \ell_j(\sigma^j(\epsilon^j))\right\}_{\mathcal{E}^j,j}$. Given a reporting strategy and an allocation, expected lifetime utility is

$$\vartheta\left(x^\sigma\right) \equiv \sum_{j=1}^{J} \sum_{\epsilon^j} \beta^{j-1} \pi\left(\epsilon^j\right) U\left(c_j\left(\sigma^j\left(\epsilon^j\right)\right), \ell_j\left(\sigma^j\left(\epsilon^j\right)\right); \epsilon_j\right).$$

The continuation value after history ϵ^j, which is denoted $V^\sigma(\epsilon^j)$, is:

$$V^\sigma\left(\epsilon^j\right) = u\left(c_j\left(\sigma^j\left(\epsilon^j\right)\right), \ell_j\left(\sigma^j\left(\epsilon^j\right)\right); \epsilon_j\right) + \beta \sum_{\epsilon_{j+1}} \pi^{j+1}\left(\epsilon_{j+1}|\epsilon_j\right) V^\sigma\left(\epsilon^{j+1}\right) \ ,$$

for all $j = 1, \ldots, J$, with $V^\sigma(\epsilon^{J+1}) = 0$. Under a truthful report strategy, the continuation value after history ϵ^j thus solves:

$$V\left(\epsilon^j\right) = u\left(c_j\left(\epsilon^j\right), \ell_j\left(\epsilon^j\right); \epsilon_j\right) + \beta \sum_{\epsilon_{j+1}} \pi^{j+1}\left(\epsilon_{j+1}|\epsilon_j\right) V\left(\epsilon^{j+1}\right),$$

for all $j = 1, \ldots, J$, with $V(\epsilon^{J+1}) = 0$.

An allocation is incentive compatible if and only if for all histories ϵ^j

$$V\left(\epsilon^j\right) \geq V^\sigma\left(\epsilon^j\right) \ ,$$

for all $\sigma \in \Sigma$. The set of incentive compatible allocations for household i is $X_{IC}(i)$. An allocation is incentive feasible if and only if the allocation for household i is incentive compatible for all households $i \in \mathcal{I}$. An allocation is feasible if and only if it is resource feasible and incentive feasible.

An allocation is efficient if and only if there is no alternative feasible allocation that makes all households weakly better off and some households strictly better off. That is, there is no alternative feasible allocation $\hat{x}$ such that $\vartheta_t\big(\hat{x}(i); \epsilon^{j-1}\big) \geq \vartheta_t\big(x(i); \epsilon^{j-1}\big)$ for all households, and $\vartheta_t\big(\hat{x}(i); \epsilon^{j-1}\big) > \vartheta_t\big(x(i); \epsilon^{j-1}\big)$ for some household. We next establish that the planning problem in the main text characterizes efficient allocations.

Proposition. Allocation x is efficient if and only if it solves the planner problem given $\vartheta_t\big(x(i); \epsilon^{j-1}\big)$ for all $i \in \mathcal{I}$ with a maximum of zero.

Proof. We show both directions by contradiction. $\Rightarrow$ If an allocation x is efficient it solves the planner problem given $\vartheta_t\big(x(t, \epsilon^{j-1}); \epsilon^{j-1}\big)$ for all $i \in \mathcal{I}$ with a maximum of zero. Suppose x does not solve the planner problem and let $\hat{x}$ denote a solution to the planner problem. Because x is feasible, the allocation $\hat{x}$ generates strictly excess resources in the first period. Construct an alternative allocation $\tilde{x}$ identical to $\hat{x}$ but increase initial consumption such that the ICs are satisfied. The allocation $\tilde{x}$ strictly Pareto dominates x, which is a contradiction.

$\Leftarrow$ If an allocation x solves the planner problem given $\vartheta_t\big(x(t, \epsilon^{j-1}); \epsilon^{j-1}\big)$ for all $i \in \mathcal{I}$ with a zero maximum, then it is efficient. Suppose that x is not efficient, then there exists an alternative feasible allocation $\hat{x}$ such that all households are better off, with some household i strictly better off. Since allocation $\hat{x}$ is feasible and delivers at least $\vartheta_t\big(x(t, \epsilon^{j-1}); \epsilon^{j-1}\big)$ for all $i \in \mathcal{I}$, $\hat{x}$ is a candidate solution to the planner problem. Construct an alternative allocation $\tilde{x}$, which is equal to $\hat{x}$ but equally reduce initial consumption for household i that is strictly better off under $\hat{x}$ (such that the ICs are satisfied). Alternative allocation $\tilde{x}$ is feasible and generates excess resources in the initial period. This contradicts that x is a solution to the planner problem.

The Lagrange function for the planning problem is separable in the allocation of each household and, therefore, we can separately characterize the solution for each household. The household planner problem is to choose a household allocation to maximize excess resources subject to the household's incentive constraints. To make this tractable, we assume only local downward incentive constraints bind at the solution. The relaxed household planner problem is formulated by replacing the set of constraints that ensure global incentive compatibility in the component planning problem, $X_{IC}(i)$, with the set of constraints that ensure the allocation satisfies all local downward incentive constraints, $X_{LD}(i)$. We then write the relaxed household planner problem recursively.

References

Benabou, Roland, 2002, "Tax and Education Policy in a Heterogeneous Agent Economy: What Levels of Redistribution Maximizes Growth and Efficiency?," *Econometrica*, 70: 481–517.

Boerma, Job, 2019, "Housing Policy Reform," Working paper, University of Minnesota.

Chamberlain, Gary, 1984, "Panel data," in: Z. Griliches and M. D. Intriligator (eds.), *Handbook of Econometrics* (Elsevier) Vol. 2, Ch. 22, pp 1247–1318.

Farhi, Emmanuel, and Ivan Werning, 2013, "Insurance and Taxation over the Life Cycle," *Review of Economic Studies*, 80(2), 596–635.

Floden, Martin, 2001, "The Effectiveness of Government Debt and Transfers as Insurance," *Journal of Monetary Economics*, 48(1): 81–108.

Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski, 2016, "Redistribution and Social Insurance," *American Economic Review*, 106(2), 359-386.

Hosseini, Roozbeh, and Ali Shourideh, 2019, "Retirement Financing: An Optimal Reform Approach," *Econometrica*, 87(4), 1205–1265.

McGrattan, Ellen, and Edward C. Prescott, 2017, "On Financing Retirement with an Aging Population," *Quantitative Economics*, 8(1), 75115.

Nishiyama, Shinichi, and Kent Smetters, 2014, "Analyzing Fiscal Policies in a Heterogeneous-Agent Overlapping-Generations Economy" *Handbook of Computational Economics*, (Elsevier) Vol. 3, Ch. 3, pp 117–160.

Stantcheva, Stephanie, 2017, "Optimal Taxation and Human Capital Policies over the Life Cycle," *Journal of Political Economy*, 125(6): 1931–1990.

FIGURE 1. PARETO EFFICIENT FRONTIER

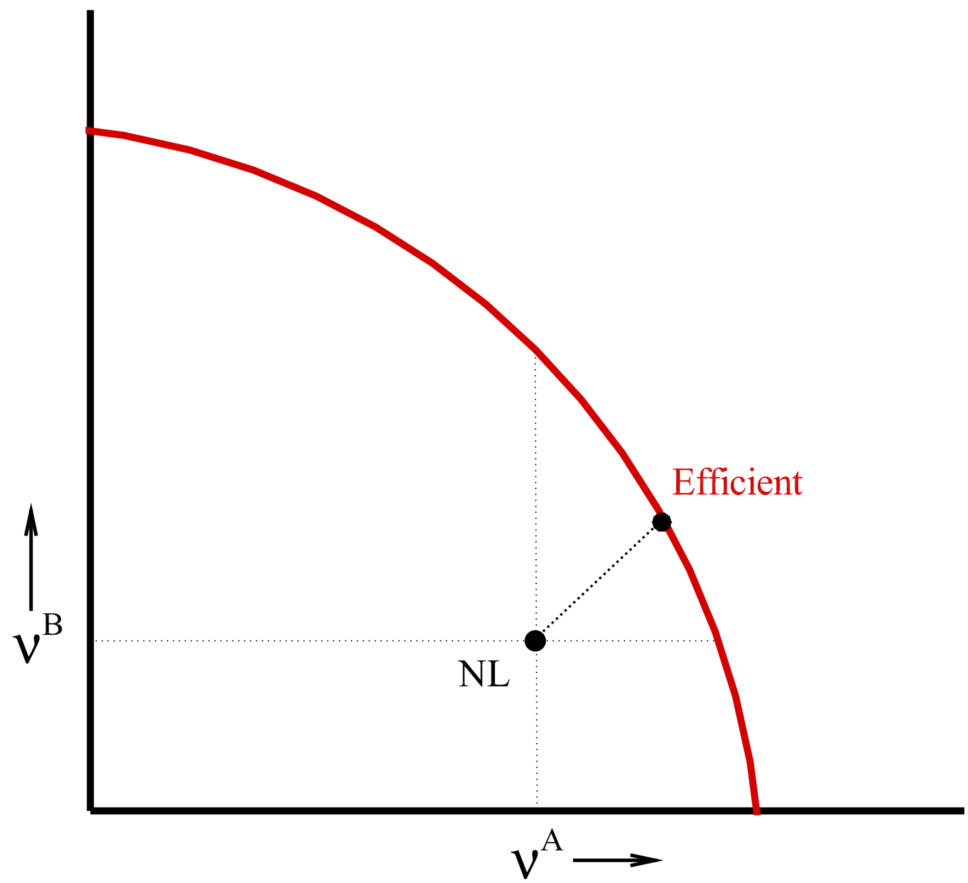

FIGURE 2. INCOME TAX SCHEDULE

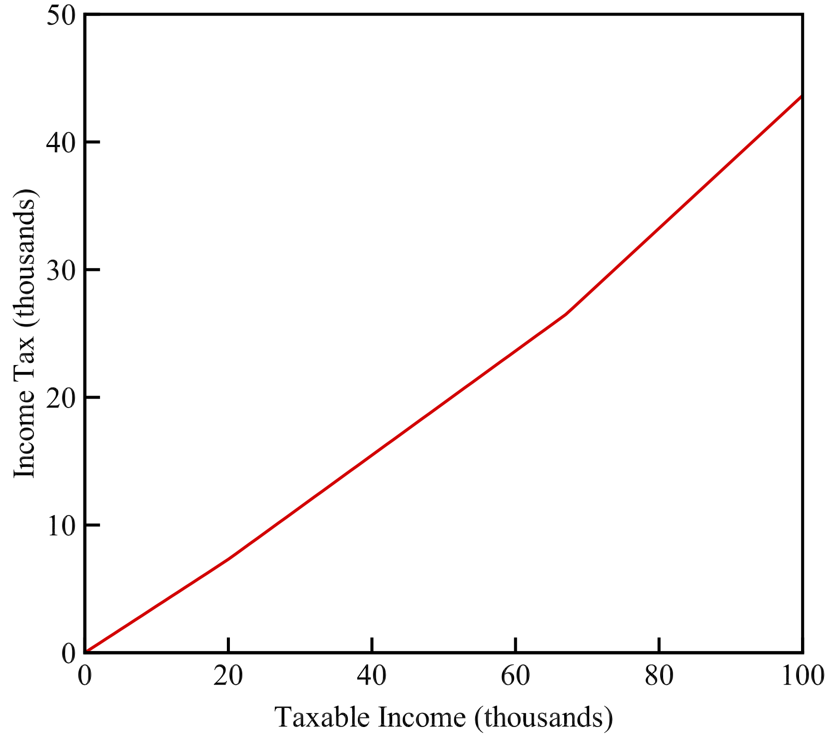

FIGURE 3. FINANCIAL ASSET TAX SCHEDULE

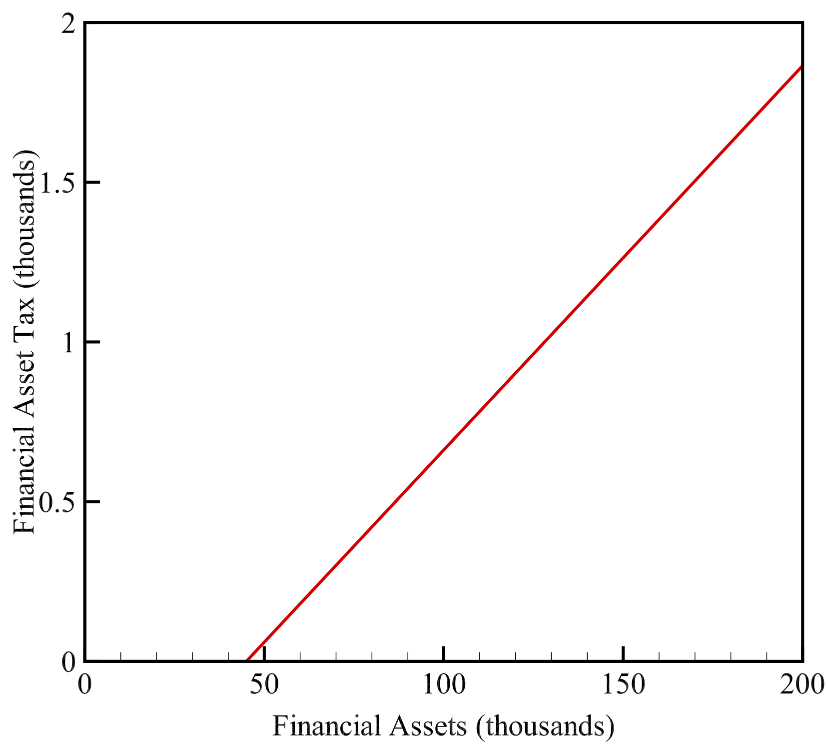

Figure 4. Wage Profiles

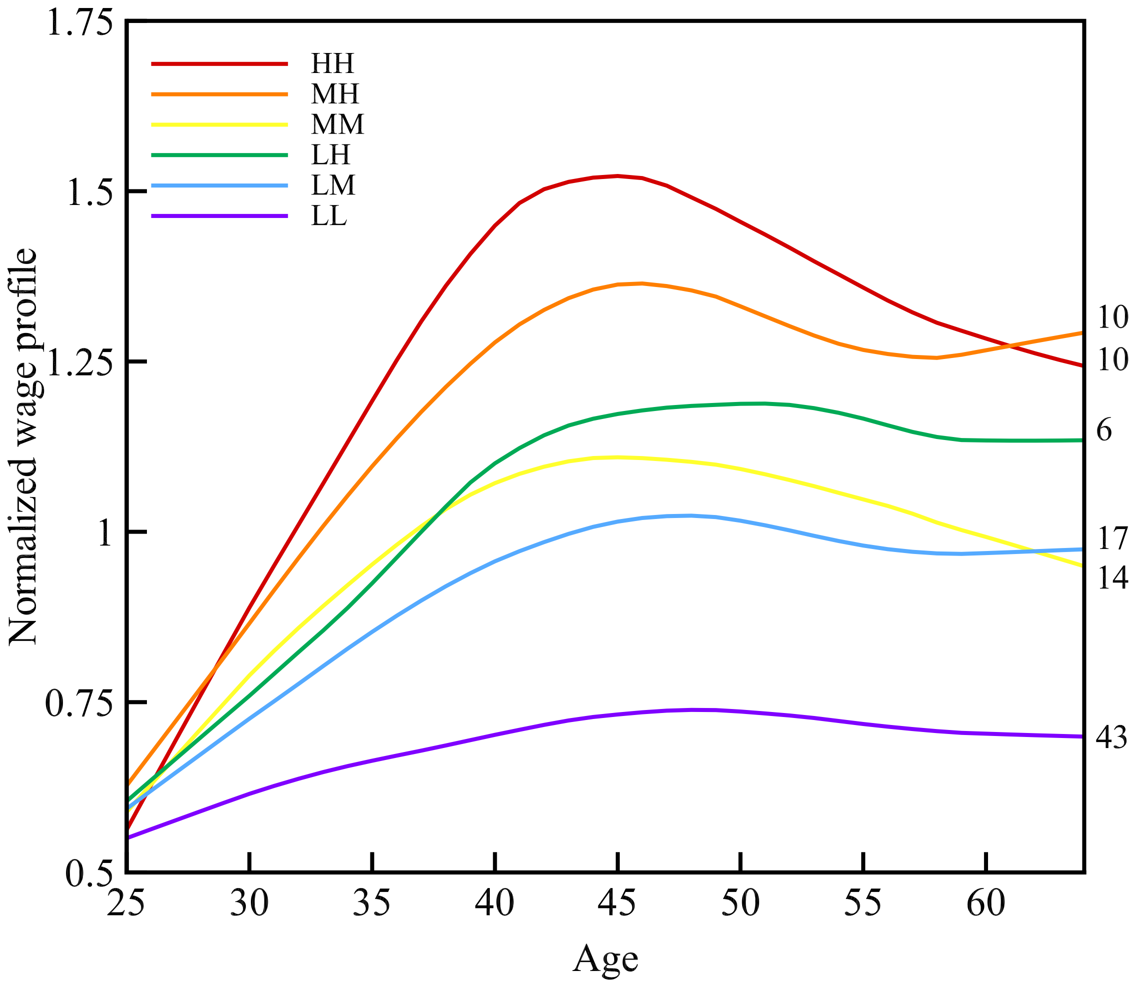

Figure 5. Labor Wedges: Baseline Model

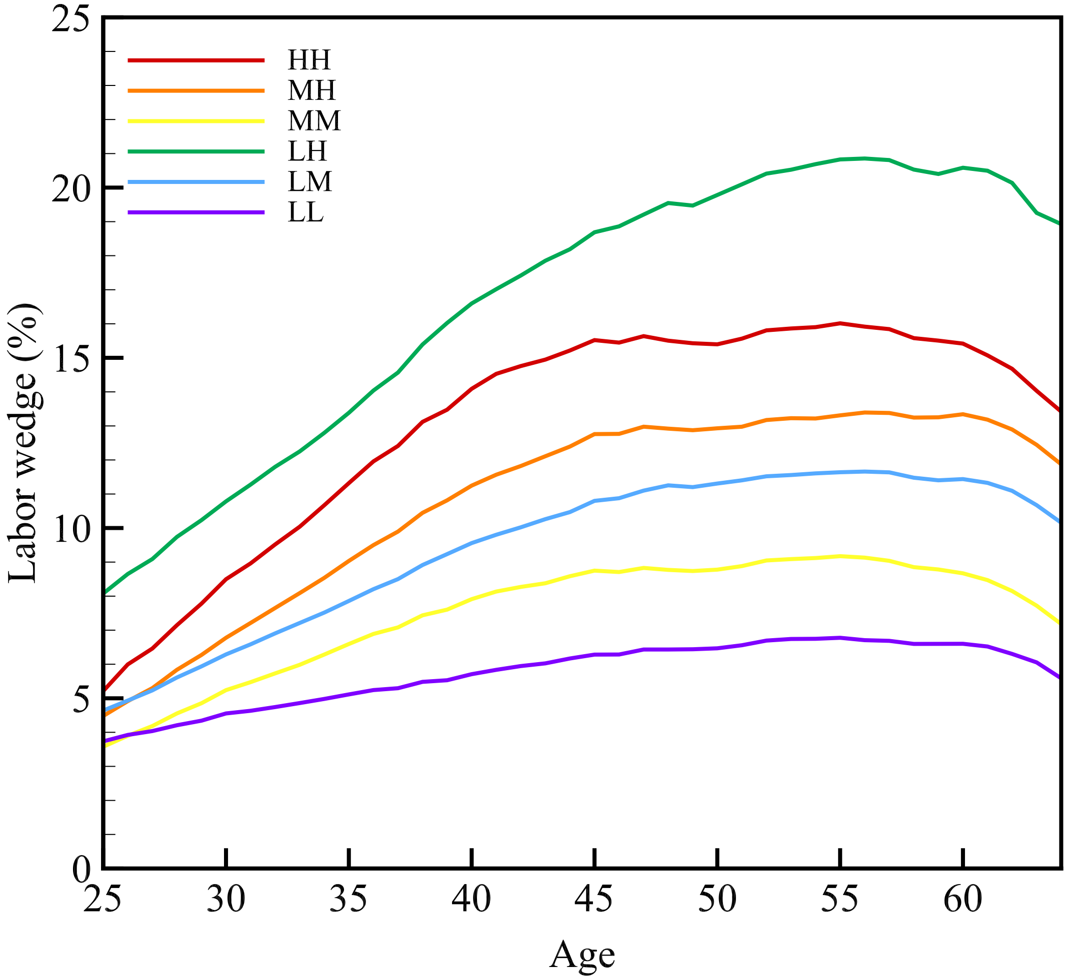

FIGURE 6. ALLOCATION LEVELS AND DISPERSION FOR LL: BASELINE MODEL

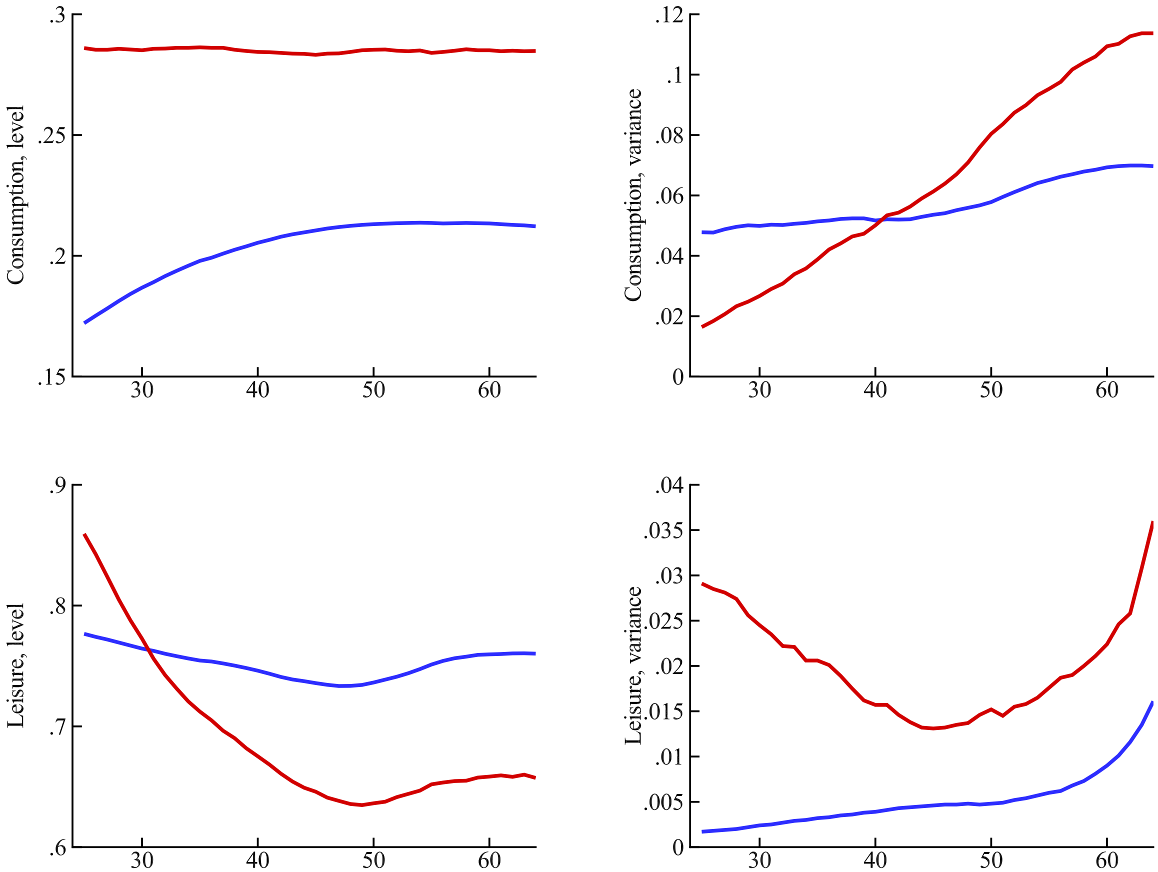

Note: Results for the positive economy are shown in blue and results for the reform problem are shown in red.

Figure 7. Allocation Levels and Dispersion for Group Low-Low

Comparison of Baseline and 40-Types Model

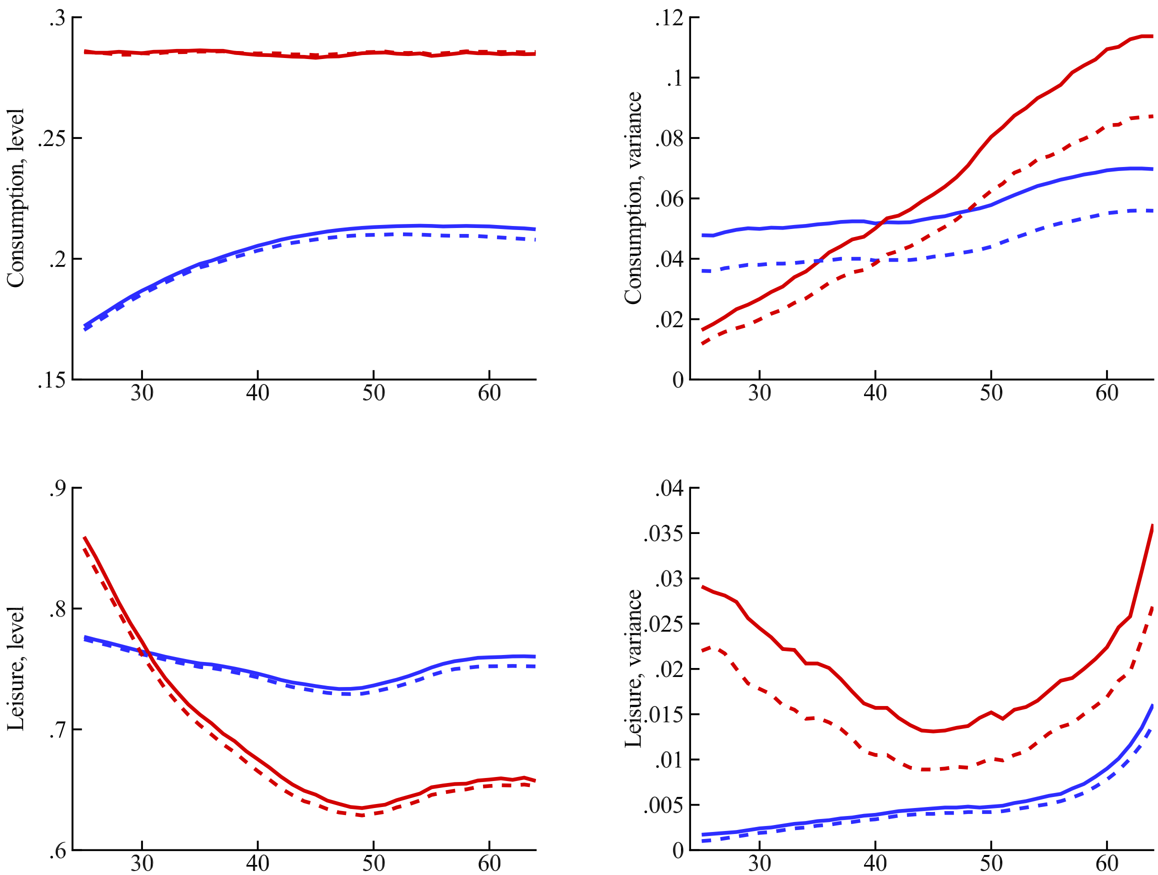

Note: Results for the positive economy are shown in blue and results for the reform problem are shown in red. The solid lines are the baseline model and the dashed lines are the 40-type model.

Figure 8. Allocation Levels and Dispersion for Group Low-Low

Comparison of Baseline and No Wage Growth Model

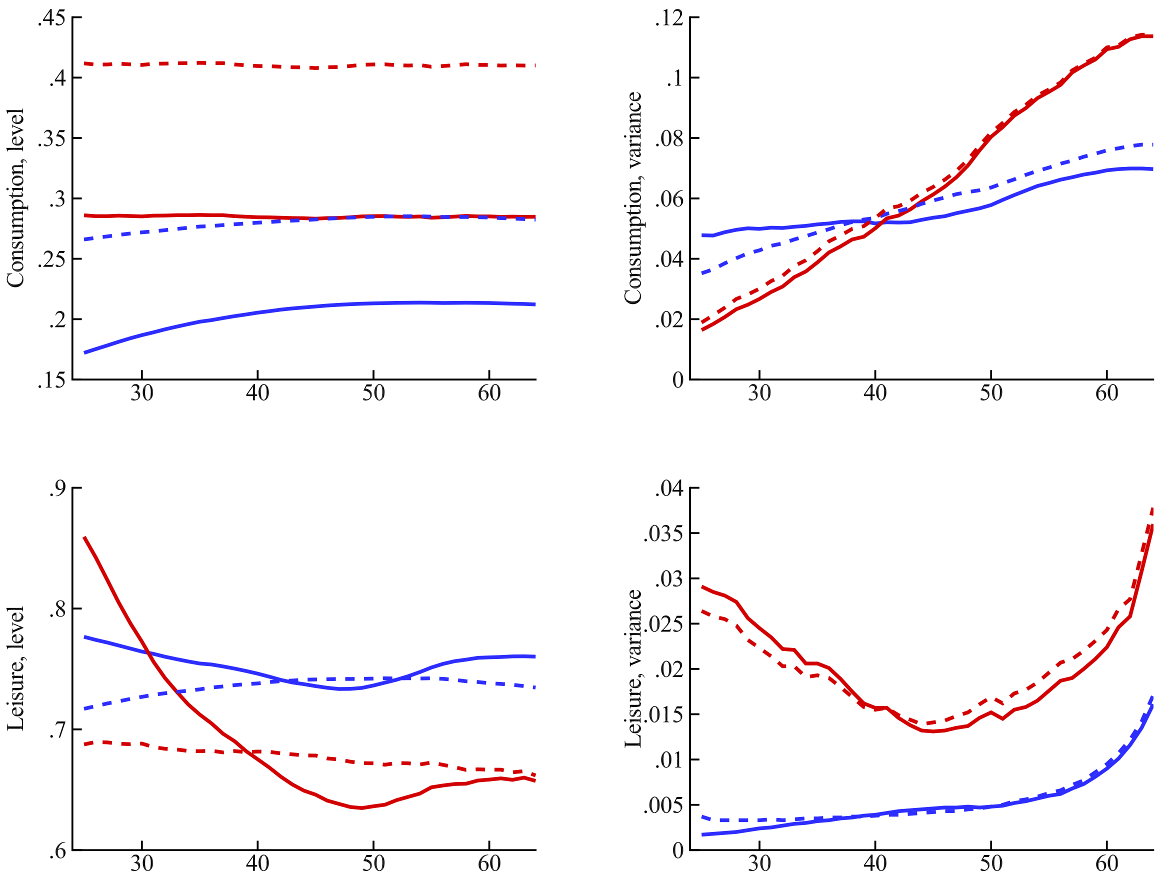

Note: Results for the positive economy are shown in blue and results for the reform problem are shown in red. The solid lines are the baseline model and the dashed lines are the no-wage-growth model.

Figure 9. Allocation Levels and Dispersion for Group Low-Low

Comparison of Baseline and Lower Shock Variance Model

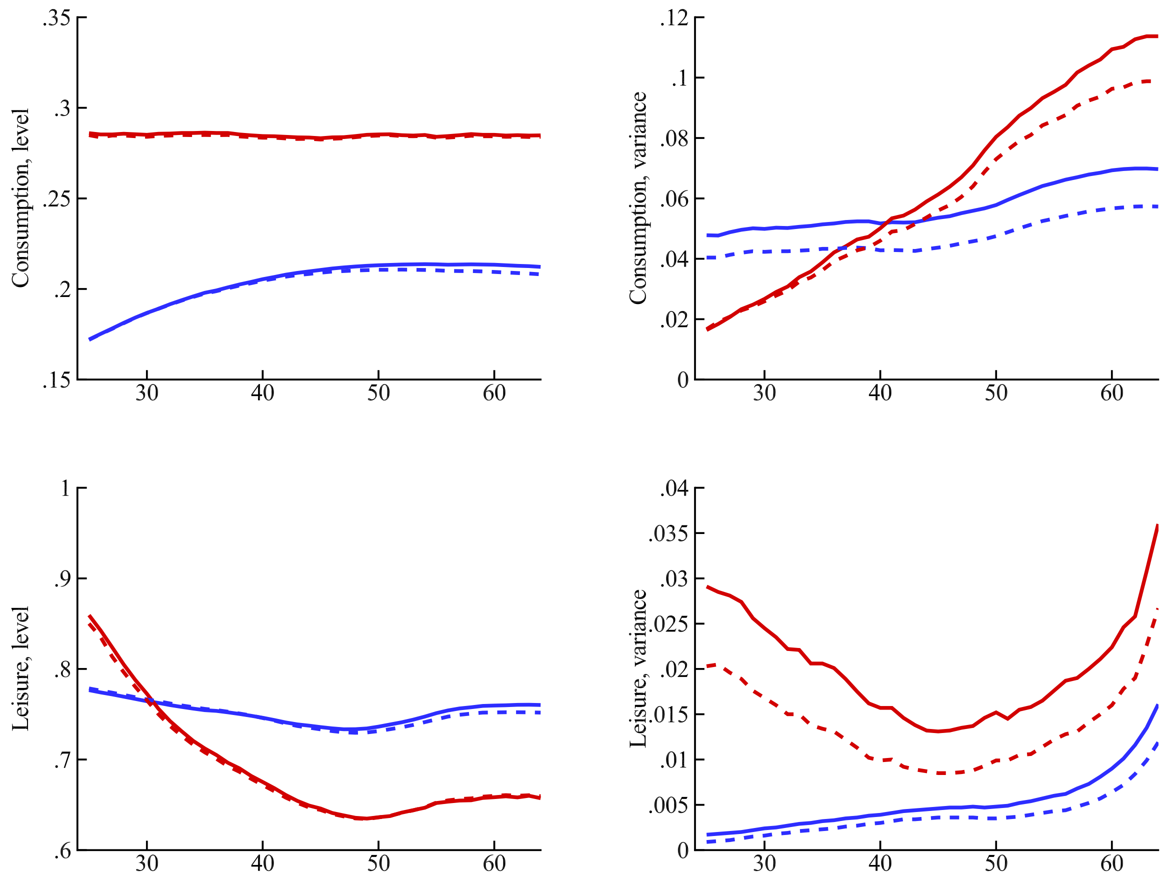

Note: Results for the positive economy are shown in blue and results for the reform problem are shown in red. The solid lines are the baseline model and the dashed lines are the lower-shock-variance model.

Table 1. Revised National Income and Product Accounts
Averages Relative to Adjusted GDP, 2000–2010

Total Adjusted Income	1.000
Labor Income	.566
Compensation of employees	.502
Wages and salary accruals	.397
Supplements to wages and salaries	.105
70% of proprietors' income	.064
Capital Income	.434
Profits	.156
30% of proprietors' income	.027
Indirect business taxes	.105
Less: Sales tax	.103
Consumption of fixed capital	.165
Consumer durable depreciation	.050
Imputed capital services	.035
Consumer durable services	.012
Government capital services	.023

See footnotes at the end of the table.

Table 1. Revised National Income and Product Accounts Averages Relative to Adjusted GDP, 2000–2010 (Cont.)

Total Adjusted Product	1.000
Consumption	.635
Personal consumption expenditures	.484
Less: Consumer durable goods	.068
Less: Imputed sales tax, nondurables and services	.088
Plus: Imputed capital services, durables	.012
Government consumption expenditures, nondefense	.222
Plus: Imputed capital services, government capital	.023
Consumer durable depreciation	.050
Tangible investment	.351
Gross private domestic investment	.177
Consumer durable goods	.068
Less: Imputed sales tax, durables	.014
Government gross investment, nondefense	.041
Net exports of goods and services	.079
Defense spending	.014

Note: The data source for national income statistics is the Dutch Bureau of Statistics. Imputed capital services are equal to 4 percent times the current-cost net stock of government fixed assets and consumer durable goods.

Table 2. Revised Fixed Asset Tables with Stocks End of Period, Averages Relative to Adjusted GDP, 2000–2010

Total Capital	5.657
Fixed assets	3.068
Businesses	1.261
Government	0.571
Households	1.236
Consumer durables	.301
Inventories	.142
Businesses	.129
Households	.013
Land	1.905
Agricultural and productive land	.420
Residential land	1.485
Oil and gas	.241

Note: The data source for national income statistics is the Dutch Bureau of Statistics.

Table 3. Household Net Worth and Government Debt Averages Relative to Adjusted GDP, 2000–2010

Household Net Worth, end of period	3.895
Assets	5.130
Tangible	2.466
Financial	2.664
Liabilities	1.236
Government Debt, end of period	.556

Note: The data source for national income statistics is the Dutch Bureau of Statistics.

Table 4. Population, Employment, and Hours Averages, 2000–2010

Population in millions	
All ages	16.3
Ages 16 to 64	10.8
Population growth (%)	
All ages	0.5
Ages 16 to 64	0.3
Annual hours per population 16-64	1,135

Note: The data source for national income statistics is the Dutch Bureau of Statistics.

Table 5. Estimated Wage Process Parameters

	Persistence		Innovation Variance	
Education Group	$\hat{\rho}$	Confidence	$\hat{\sigma}_u^2$	Confidence
Low, Low	.9542	(.9515,.9575)	.0096	(.0093,.0102)
Low, Medium	.9660	(.9610,.9692)	.0087	(.0083,.0096)
Low, High	.9673	(.9628,.9710)	.0162	(.0153,.0176)
Medium, Medium	.9570	(.9536,.9612)	.0099	(.0091,.0103)
Medium, High	.9616	(.9520,.9782)	.0109	(.0082,.0124)
High, High	.9564	(.9501,.9582)	.0172	(.0164,.0184)

Table 6. Welfare Gain Decomposition: Baseline Model
Total Welfare Gain of 20%

	Consumption		Leisure	
Education group	Δ_c^L	Δ_c^D	Δ_ℓ^L	Δ_ℓ^D
Low, Low	27	2	-8	-2
Low, Medium	25	4	-9	-1
Low, High	20	11	-8	-2
Medium, Medium	27	4	-10	-1
Medium, High	25	7	-11	-1
High, High	21	17	-14	-5

Table 7. Welfare Gain Decomposition: 40 Types Model
Total Welfare Gain of 20%

	Consumption		Leisure	
Education group	Δ_c^L	Δ_c^D	Δ_ℓ^L	Δ_ℓ^D
Low, Low	28	1	-8	-2
Low, Medium	28	4	-11	-1
Low, High	28	10	-16	-2
Medium, Medium	29	4	-11	-2
Medium, High	29	7	-14	-2
High, High	33	14	-22	-5

Table 8. Welfare Gain Decomposition: No Growth Model
Total Welfare Gain of 18%

	Consumption		Leisure	
Education group	Δ_c^L	Δ_c^D	Δ_ℓ^L	Δ_ℓ^D
Low, Low	28	0	−7	−4
Low, Medium	27	0	−6	−3
Low, High	25	0	−5	−2
Medium, Medium	28	0	−6	−4
Medium, High	27	0	−6	−3
High, High	25	1	−6	−3

Table 9. Welfare Gain Decomposition: Lower Variance Model
Total Welfare Gain of 18%

	Consumption		Leisure	
Education group	Δ_c^L	Δ_c^D	Δ_ℓ^L	Δ_ℓ^D
Low, Low	25	1	−7	−2
Low, Medium	25	3	−9	−1
Low, High	24	7	−12	−1
Medium, Medium	25	3	−8	−1
Medium, High	24	5	−10	−1
High, High	23	10	−12	−2