Optimal Social Security Reforms: a Bayesian Quantitative Analysis

Marco Cozzi, Queen's University

PRELIMINARY AND INCOMPLETE, PLEASE DO NOT CITE

February 2014

Abstract

This paper characterizes numerically the optimal pension benefits scheme when two competing specifications for the stochastic income process govern the extent of labor market risk. The analysis embeds elements of both Bayesian empirical methods and robust social security reform. A quantitative Macroeconomic model with heterogeneous agents is exploited to address both which stochastic income process is more likely to have generated the data, and to compute the optimal benefits function for the social security system in an economy with rich labor income dynamics. Under the Minimal Econometric Interpretation of the economic model, proposed by Geweke (2010), the predictive distributions of a set of population moments for the two models are compared to their empirical counterparts, through an a-theoretical Bayesian Panel VAR with implications for the same moments: the autocorrelations of earnings and hours worked. Preliminary results seem to give more weight to the model of wage dynamics with slope heterogeneity, but not overwhelmingly so, as the Bayes Factor is approximately $3:1$. Moreover, the optimal replacement rate of the social security system is found to depend heavily on the nature of labor income risk. Finally, the optimal replacement rate with either specification of wage risk is found to be higher than what is typically obtained in similar studies.

JEL Classification Codes: C52, C54, C68, D52, E21, E62, H55.

Keywords: Social Security, Heterogeneous Agents, Incomplete Markets, Ex-ante Policy Evaluation, Welfare, Bayesian Methods.

Acknowledgements: *I am grateful to seminar participants at Guelph University for useful comments and suggestions, and to SHHRC for providing financial support for this project (IDG Grant # 430-2013-511).*
Contact details: Department of Economics, Queen's University, 94 University Avenue, Kingston, ON K7L 3N6, Canada. Tel: 1-613-533-2264, Fax: 1-613-533-6668, Email: mcozzi@econ.queensu.ca.

1 Introduction

The extent of labor market risk is central to both many economic decisions and to the design of several public policies. In particular, the contributions towards a public pension system depend crucially on the evolution of labor income earnings, which is uncertain, because of uninsurable labor market risk. Figure (1) presents a fairly popular plot: the (residual) variances of log-wages and log-earnings in the PSID data, for the 1968-1996 period. The main message is that, during the life-cycle, there is a spectacular increase in wage inequality, which is not due to observable factors.

[Figure 1 about here]

The increase in the variance of log-wages is almost linear, and it is pretty uncontroversial in the literature. Differently, the increase in the variance of log-earnings displays a more complex pattern, which in turn depends on the features of the econometric model, such as the specification of the first stage regression.

Two main views have emerged in the literature to explain the rise in labor income inequality over the life-cycle. One view, proposed among others by Storesletten, Telmer and Yaron (2004b) and Heathcote, Storesletten and Violante (2010), posits that upon labor market entry workers are relatively homogeneous. However, they face a high degree of uncertainty over their working life, because they experience labor income shocks that are highly persistent. The effects of these shocks compound over time, leading to an increasing inequality in the labor market outcomes over the life-cycle, as analyzed by Storesletten, Telmer and Yaron (2004a). Hereafter I will refer to this specification as the Standard Income Profiles (SIP).

The other view, studied by Guvenen (2007), Guvenen (2009) and Huggett, Ventura, and Yaron (2011), posits that workers markedly differ in their earning capabilities over the life-cycle already at the beginning of their working life. According to this view, inequality is primarily generated by this channel (namely, the heterogeneity in the wage growth), because workers experience labor income shocks that are only mildly persistent. I will refer to this specification as the Heterogeneous Income Profiles (HIP).

The typical strategy used to estimate these models is to specify an exogenous process for the wage dynamics, consider the theoretical covariances implied by the statistical model, and use panel data to estimate the parameters with a minimum distance estimator, as discussed by Heathcote, Storesletten and Violante (2009). Unfortunately, Guvenen (2009) shows that the power of the tests used to examine if the slope heterogeneity is part of the Data Generating Process (hence differentiating between the two models) is really low, making the inference very unreliable.

A first contribution of this paper is to provide further evidence on the nature of labor market risk. My approach is going to be different from most contributions in the literature, as I will use the economic model to draw inferences on which statistical model of wages (between HIP and SIP) is more likely to have generated the data. In a nutshell, I will posit empirically motivated prior distributions for the parameters for the two statistical models for *wages*, and I will exploit the economic model to consider the implications for two (observable) endogenous outcomes: the autocorrelation of labor *earnings* and the autocorrelation of hours worked. Geweke

(2010) labels this procedure the Minimal Econometric Interpretation of the economic model. One of the main advantages of this simple Bayesian method is allowing to compare two models that can be both misspecified, because it doesn't rely on tests that must assume that one of the models is the true DGP. This approach seems particularly suited for macroeconomics research, where stylized models are still extensively used for both quantitative and welfare analyses. Furthermore, under the Minimal Econometric Interpretation, the predictions of the economic model can be mapped to the data without the need to either maximize the likelihood function or compute the posteriors, steps that are still remarkably expensive for this class of models, where agents' heterogeneity is a key ingredient. In this framework, the predictive distributions of a set of population moments for the two models are compared to their empirical counterparts, through an a-theoretical Bayesian Panel VAR with implications for the same moments. This step is necessary, as the economic models are interpreted to inform us about population moments, which are by definition unobservable: linking the model to the data is then accomplished with an exogenous econometric model. The model comparison between SIP and HIP is performed by considering the marginal likelihoods: preliminary results seem to give more weight to the model of wage dynamics with slope heterogeneity, but not overwhelmingly so, as the Bayes Factor is approximately $3 : 1$.

In the first stage of the analysis the model will consider a specification of the pension system that mimics the actual one in place in the US economy. In particular, the pension benefits will be specified according to a close approximation of the schedule defining the Primary Insurance Amount (PIA).

Irrespective of whether SIP or HIP is assumed, in both cases workers cannot easily self-insure with their saving behavior to achieve a smooth consumption profile. As standard in models working with a similar set-up, such as Imrohoroglu, Imrohoroglu, and Joines (1998) or Fuster, Imrohoroglu and Imrohoroglu (2003), the presence of market incompleteness and borrowing constraints gives room to public intervention, because appropriate redistribution schemes can be welfare improving.

I focus the attention on a reform of the pension system, considering a government whose objective is to maximize welfare (or the share of agents whose welfare is positively affected by the policy change) by choosing the replacement rates of the public pensions. Finally, the contributions to the system will be set residually, in order for the pension system to be self-financing.

I characterize the optimal benefit scheme (in a way that will be made precise in Section 6) under either SIP or HIP. At the same time, the approach embeds a global robustness analysis, as the two models are solved with a large number of calibrations.

In this set-up the nature of the labor income shocks is crucial for the quantitative findings. Different individuals experience different labor market histories, which lead to different contributions towards the pension system, and to different incomes that the retirees will receive upon leaving the labor force. Intuitively, the higher the variance and the persistence of the income shocks, the more progressive the pension system should be, because risk averse individuals demand a substantial redistribution to compensate for the bad labor market outcomes.

As for the empirical implementation of the models, I begin by consider a Bayesian Prior Analysis, namely I specify prior distributions for the models' parameters, in particular the ones describing the two exogenous stochastic processes for wages. A sensible way for selecting these priors is to consider uniform distributions over the 95% confidence intervals obtained from the reduced form estimates of both the SIP and HIP exogenous stochastic processes.

With a prior assumption on the probability of the two wage processes, I can solve the economy under a new proposal for the pension system. I can do so for many different realizations of the model's parameters to compute the average welfare implied by this policy reform. By repeating this procedure for a fine grid over the parameters mapping into the replacement rates, I am able to characterize the pension system reform that achieves the highest welfare, which at the same time is robust to uncertainty about the nature of labor market risk. This robustness has two different layers. First, for a given view of labor market risk, it allows for parameter uncertainty stemming from the intrinsic uncertainty of the reduced form estimators, by considering the implications of different values for the parameters representing the variances and correlations of the labor market shocks. Second, at a deeper level, it is robust to the government's lack of knowledge on the true nature of labor market risk. Crucially, the government is uncertain about which stochastic mechanism, SIP vs. HIP, is driving the labor market risk, and it will have to confront this lack of knowledge when evaluating the welfare associated with a contemplated policy change.

Although already very rich (and computationally complex), a Bayesian Prior Analysis is somewhat unsatisfactory. It does provide a careful description of the likelihood of the empirical outcomes for a given set of parameters' priors, but it does not confront the results with the information available in the data. Hence, I will also consider the Minimal Econometric interpretation of the DSGE model. More precisely, under this view the model conveys information only on a set of population moments, such as the autocorrelations of the hours worked or of labor earnings, which are unobservable. The Panel Study of Income Dynamics can be easily used to specify a simple Bayesian dynamic panel data model, with implications for the same moments, achieving the goal of relating the population moments to the sample ones. The attractiveness of this approach is that, even in the presence of misspecified economic models, it offers evidence on which model provides a better description of the data. Since there are two main competing formulations of the nature of labor market risk, SIP and HIP, this analysis could represent an important contribution to shed more light on which explanation fits the data better, and how we should tackle its consequences. The main contribution of this additional analysis is to give the prior probabilities for the two competing models of labor market risk in the government's robust policy problem a sound empirical foundation.

The rest of the paper is organized as follows. Sections 2 and 3 briefly review the two statistical models of wage dynamics, and the current US pension system, respectively. Section 4 describes the model, and the equilibrium. Section 5 discusses the calibration and how the models are brought to the data. Section 6 presents the counterfactual analysis and the related results. Section 7 concludes. Some appendices are also included: they discuss the details of the numerical methods.

2 Two Models of Wage Dynamics: SIP and HIP

Before presenting the economic model, it is useful to discuss the two processes for wage dynamics. These are two competing formulations of the nature of labor market risk, SIP and HIP: the first part of my analysis sheds more light on which explanation fits the data better, and the second one how we should tackle its consequences, in terms of the design of an optimal pension system. An additional contribution of the analysis is to give the prior probabilities for the two competing models of labor market risk in the government's robust policy problem a sound empirical foundation.

In terms of notation, ϵ_{it} stands for the total efficiency units of labor an agent is endowed with in period t of his working life, α_i is a time invariant fixed effect, which affects permanently the level of the worker's productivity, ν_{it} is a purely transitory shock, and ε_{it} is a serially correlated shock, with an AR(1) specification.

Under SIP, the stochastic process for the efficiency units of labor is represented by:

$$\begin{cases} \epsilon_{it} = \alpha_i + \varepsilon_{it} + \nu_{it} \\ \varepsilon_{it} = \rho_\varepsilon \varepsilon_{it-1} + \eta_{it} \\ \alpha_i \overset{iid}{\sim} N(0, \sigma_\alpha^2), \nu_{it} \overset{iid}{\sim} N(0, \sigma_\nu^2), \eta_{it} \overset{iid}{\sim} N(0, \sigma_\eta^2) \end{cases}$$

Differently, under HIP a second time invariant fixed effect β_i affects the growth rate of the worker's productivity, and the stochastic process for the efficiency units of labor is represented by:

$$\begin{cases} \epsilon_{it} = \alpha_i + \beta_i t + \varepsilon_{it} + \nu_{it} \\ \varepsilon_{it} = \rho_\varepsilon \varepsilon_{it-1} + \eta_{it} \\ (\alpha_i, \beta_i) \overset{iid}{\sim} N(0, 0; \sigma_\alpha^2, \sigma_\beta^2, \sigma_{\alpha\beta}), \nu_{it} \overset{iid}{\sim} N(0, \sigma_\nu^2), \eta_{it} \overset{iid}{\sim} N(0, \sigma_\eta^2) \end{cases}$$

Although HIP nests SIP, the available empirical estimates for the two processes on PSID data lead to some noticeable differences.[1]

[Table 1 about here]

As shown in Table 1, the estimates by Guvenen (2009) highlight how the autocorrelation of the persistent shocks ρ_ε is drastically different between the two specifications: $\widehat{\rho}_{\varepsilon,SIP} = 0.988$ and $\widehat{\rho}_{\varepsilon,HIP} = 0.821$. Under SIP, highly persistent shocks that are almost a unit root are needed to account for the increasing variance of log-wages, while a substantial heterogeneity in the wage growth parameter β_i achieves this goal under HIP. It goes without saying that this discrepancy implies that these shocks are potentially very hard to insure against under SIP, can lead to very different behaviors in terms of both savings and labor supply, and can have major implications for the design of public policy, as they greatly affect the demand for redistribution.

[1] Since the transitory shocks ν_{it} are typically easy to insure against, they are going to be omitted from the economic model, allowing to consider one less continuous state variable in the dynamic programming problem.

3 The US Pension System

The current pension system in place in the US is too complicated for a quantitative analysis. The relevant statistic to determine the benefits retirees are entitled to are the 35 highest labor earnings during a person's working life. This is computationally intractable, as it would require 35 continuous state variables: I will follow Imrohoroglu and Kitao (2012) to simplify the problem. The relevant statistic for the PIA is going to be the average labor earnings over the life cycle. This simplification makes the analysis feasible, introducing just a relatively small error, as the typical working life is less than 40 years.

In 2008 the actual formula was as follows, where y stands for the individuals' average labor earnings:

$$PIA = \begin{cases} 0.9 * y, & \text{if } y < \$8,532 \\ \$7,679 + 0.32 * (y - \$8,532), & \text{if } \$8,532 \leq y < \$51,456 \\ \$21,414 + 0.15 * (y - \$51,456), & \text{if } \$51,456 \leq y < \$102,000 \\ \$28,995, & \text{if } y \geq \$102,000 \end{cases}$$

Figure (2) depicts the function describing the PIA. It describes the benefits pensioners are entitled to as a function of their average labor earnings, which are a measure of their overall contributions to the pension system.

[Figure 2 and 3 about here]

Figure (3) provides a similar plot, but on the x-axis there are the relative contributions, namely the overall contributions divided by the average ones (PIAR). In the quantitative analysis I will focus on this specification, because it allows for an interesting trade-off between agents with different accumulated contributions that will be discussed below.

The characterization of the optimal benefit scheme requires to take a stand on the family of functions that represents the possible PIAR's. As seen from the formula above, the key features of the PIAR are four slopes and three thresholds. However, optimizing over these seven variables is not desirable. Following the idea of Gouveia and Strauss (1994), applied by Conesa, Kitao and Krueger (2009) in their study on the optimal capital income tax, I will use a parsimonious yet very flexible specification for the PIAR. The formula will be a tractable but general alternative: $PIAR = p_0 + p_1 * \left(\frac{y}{\bar{y}}\right)^{p_2}$. Fitting this function with non-linear least squares delivers the following estimates for the parameters, that describe the benefits in the current US social security system. $PIAR = \widehat{p}_0 + \widehat{p}_1 * \left(\frac{y}{\bar{y}}\right)^{\widehat{p}_2} = -0.00984 + 0.538 * \left(\frac{y}{\bar{y}}\right)^{0.455}$.

[Figure 4 about here]

Figure 4 plots both the fitted and the actual PIA schedules. As it is apparent from the figure, there are only minor discrepancies between the actual benefits pensioners are entitled to, and the ones prescribed by the approximation. Furthermore, since the estimate for $\widehat{p}_0 \approx 0$, this suggests to neglect it, reducing the number of variables in the optimization problem from seven to two.

In the first step of the analysis, the benchmark economy will be calibrated with the actual pension scheme, i.e. with the pair of values ($\widehat{p}_1 = 0.538, \widehat{p}_2 = 0.455$) obtained in the approximation. Then, in the counterfactual analysis I will look for the optimal PIA system, i.e. I will search for the pair (p_1^*, p_2^*) that achieves the highest average welfare $W\left(p_1^*, p_2^*\right)$.[2]

4 The Model Economy

The model is an extension of Huggett (1996). Time is discrete. The economy is populated by a measure one of finitely lived agents facing an age-dependent death probability π_t^d. Age is denoted with t and there are T overlapping generations, each consisting of a continuum of agents. At age T_R all alive agents retire. The population grows at rate g_n.[3]

Preferences: Agents' preferences are assumed to be represented by a time separable utility function $U(.)$. Agents' utility is defined over stochastic consumption and leisure sequences $\{c_t, l_t\}_{t=1}^T$: their aim is to choose how much to consume (c_t), how much to save in an interest bearing asset (a_{t+1}), and how many hours $h_t = 1 - l_t$ to work in each period of their lives, in order to maximize their objective function. The agents' problem can be defined as:

$$\max_{\{c_t, a_{t+1}, l_t, h_t\}_{t=1}^T} E_0 U(c_0, l_0, c_1, l_1, ...) = \max_{\{c_t, a_{t+1}, l_t, h_t\}_{t=1}^T} E_0 \sum_{t=1}^T \delta^{t-1} \left[\prod_{j=1}^t \left(1 - \pi_j^d\right) \right] u(c_t, l_t)$$

where E_0 represents the expectation operator over the idiosyncratic sequences of shocks, and $\delta > 0$ is the subjective discount factor. I assume that $u(c_t, 1 - h_t) = \frac{\left[c_t^{1-\lambda}(1-h_t)^{\lambda}\right]^{1-\gamma} - 1}{1-\gamma}$, that is the per-period utility function is strictly increasing, strictly concave, satisfies the Inada conditions, and has a EIS$= 1/\gamma$.

Endowments: Agents differ in their labor endowments $\epsilon_{t,\varepsilon,f}$. There are several channels that contribute to the determination of the total efficiency units that the workers supply in the labor market. First, there is a deterministic age component e_t, which is the same for all agents. Second, there is a stochastic component ε, whose log follows a stationary $AR(1)$ process: $\log \varepsilon_t = \rho_\varepsilon \varepsilon_{t-1} + \eta_t$, with $\eta_t \sim N(0, \sigma_\eta^2)$. Third, there is a fixed effect component f, whose realizations are drawn from either a normal (under SIP, with $\alpha \sim N(0, \sigma_\alpha^2)$), or a bivariate correlated normal (under HIP, with $(\alpha_i, \beta_i) \sim N(0, 0; \sigma_\alpha^2, \sigma_\beta^2, \sigma_{\alpha\beta})$). The total efficiency units a worker is endowed with are the product of these components, and total labor earnings are $wh_t\epsilon_{t,\varepsilon,f} = w \times h_t \times e_t \times \varepsilon \times \alpha$ under SIP and $wh_t\epsilon_{t,\varepsilon,f} = w \times h_t \times e_t \times \varepsilon \times \alpha \times \beta$ under HIP.

After the common retirement age T_R, the labor endowment drops to zero, and the agents receive a pension $p\left(y\right)$ paid for with the contributions of the economically active agents. The pension is a function of the average

[2] There are other alternatives to my choice, such as maximizing the welfare of the new-borns or the probability of being voted against the status-quo $(\widehat{p}_1, \widehat{p}_2)$. Moreveor, all the counterfactual schemes will satisfy a balanced budget, and public debt is not going to be allowed to be issued to finance the pensions of the retirees.

[3] The numerical methods and the computational algorithms are reported in Appendices A and B.

lifetime labor earnings y, and agents pay proportional taxes (τ_p) to contribute to the budget-balanced pension scheme. They also finance the public expenditure G and a public transfer TR with proportional capital (τ_a) and labor/pension earnings (τ_w). Agents can insure against their mortality risk in a perfectly competitive market for annuities. As a consequence, on average agents die with zero wealth and there are no accidental bequests. All agents enter the economy with a zero asset endowment and with the average realization of the stochastic component of labor earnings, which is normalized to 1.

Technology: The production side of the economy is modeled as a constant returns to scale technology of the Cobb-Douglas form, which relies on aggregate capital K and labor L to produce the final output Y.

$$Y = F(K, L) = K^{\phi} L^{1-\phi}.$$

Capital depreciates at the exogenous rate d and firms hire capital and labor every period from competitive markets. The first order conditions of the firm give the expressions for the net real return to capital r and the wage rate w:

$$r = \phi \left(\frac{L}{K} \right)^{1-\phi} - d, \tag{1}$$

$$w = (1 - \phi) \left(\frac{K}{L} \right)^{\phi}. \tag{2}$$

Government: The government carries out some public expenditure G and rebates some lump-sum transfers TR to every household. In order to finance the cost of these purchases and subsidies, both capital (τ_a) and labor (τ_w) taxes are levied. G, TR and τ_a are going to be policy parameters set exogenously, while labor taxes are set residually to guarantee a balanced budget.

Other market arrangements: The final good market is competitive, and firms hire capital and labor every period from competitive markets. Capital is supplied by rental firms that borrow from the agents at the risk-free rate r and invest in physical capital. There are no state-contingent markets to insure against the income risk, but workers can self-insure by saving into the risk-free asset. The agents also face an exogenous borrowing limit, denoted as b, which is set to $b = 0$ to avoid capital income taxes turning into subsidies for agents in debt.

4.1 Recursive Representation

In this Section, first the problem of the agents in their recursive representation is defined, then I provide a formal definition of the equilibrium concept used in this model, the recursive competitive equilibrium. The individual state variables are: age $t \in \mathcal{T} = \{1, ..., T\}$, the fixed effect $f \in \mathcal{F} = \{\alpha_{\min}, ..., \alpha_{\max}\} \times \{\beta_{\min}, ..., \beta_{\max}\}$, the persistent labor endowment shock $\varepsilon \in \mathcal{E} = \{\varepsilon_{\min}, ..., \overline{\varepsilon}, ..., \varepsilon_{\max}\}$, asset holdings $a \in \mathcal{A} = [-b, \overline{a}]$, and past average labor earnings $y \in \mathcal{Y} = [\underline{y}, \overline{y}]$. Notice that ε is discretized with the Rouwenhorst method, using an 11-state Markov chain. The transition function of the labor endowment shocks is represented by the matrix

$\Pi(\varepsilon', \varepsilon) = [\pi(i,j)]$, where each element $\pi(i,j)$ is defined as $\pi(i,j) = \Pr\{\varepsilon_{t+1} = i | \varepsilon_t = j\}$, $i, j \in \mathcal{E}$. In every period the total efficiency units are given by $\epsilon_{t,\varepsilon,f} = e_t \varepsilon f$. The stationary distribution of working-age agents is denoted by $\mu_t(a, y, \varepsilon, f)$ while that of retirees with $\mu_t^R(a, y)$. Φ_t denotes the share of each cohort t in the total population. These satisfy the recursion $\Phi_{t+1} = \left(\frac{1-\pi_j^d}{1+g_n}\right)\Phi_t$, and are normalized to add up to 1.

Problem of the agents: The model is solved backwards, starting from the terminal age T and with the assumption that the terminal utility value is zero, i.e. $V_{T+1} = 0$.

Problem of the retirees: The value function of an age-t retired agent whose current asset holdings are equal to a is denoted with $V_t^R(a, y)$. The problem of these agents can be represented as follows:

$$V_t^R(a, y) = \max_{c, a'} \left\{u(c, 1) + \delta\left(1 - \pi_t^d\right) V_{t+1}^R(a', y')\right\}, t \geq T_R \tag{3}$$

$$s.t.$$

$$c + a' = \left(\frac{1 + (1 - \tau_a) r}{1 - \pi_t^d}\right) a + (1 - \tau_w) p(y) + tr$$

$$p(y) = p_1 \left(\frac{y}{\overline{y}}\right)^{p_2}$$

$$c \geq 0, \quad a' > 0, \quad y' = y, \quad \epsilon_{t,\varepsilon,f} = 0$$

In the budget constraint notice the presence of the pension payment $p(y) = p_1 (y/\overline{y})^{p_2}$, and that the interest rate is adjusted by both the capital tax τ_a and by $\left(1 - \pi_j^d\right)$, because of the assumption on the availability of annuities. Retired agents pay proportional pension income taxes at rate τ_w, capital income taxes at rate τ_a, and receive a lump-sum transfer tr. They are no longer productive, because their efficiency units endowment is equal to zero, and they have stopped contributing to the pension system, hence their average labor income is constant over time.

Problem of the working-age agents - SIP and HIP: The value function of a working-age agent whose current asset holdings are equal to a, whose cumulated average labor income is y, whose current efficiency units shock is ε and whose time-constant fixed effect is f is denoted with $V_t(a, y, \varepsilon, f)$.

In the SIP model, the problem of these agents can be represented as follows:

$$V_t(a,y,\varepsilon,\alpha) = \max_{c,a',h} \left\{ u(c,1-h) + \delta\left(1-\pi_t^d\right) \sum_{\varepsilon'} \pi(\varepsilon',\varepsilon) V_{t+1}(a',y',\varepsilon',\alpha) \right\} \tag{4}$$

$$s.t.$$

$$c + a' = \left(\frac{1+(1-\tau_a)r}{1-\pi_t^d}\right) a + (1-\tau_w-\tau_p)\, wh\epsilon_{t,\varepsilon,\alpha} + tr$$

$$y' = [(t-1)y + wh\epsilon_{t,\varepsilon,\alpha}]/t$$

$$a_0 = 0, \quad c \geq 0, \quad a' > -b, \quad 1 \geq h \geq 0$$

In the HIP model, the problem of these agents can be represented as follows:

$$V_t(a,y,\varepsilon,\alpha,\beta) = \max_{c,a',h} \left\{ u(c,1-h) + \delta\left(1-\pi_t^d\right) \sum_{\varepsilon'} \pi(\varepsilon',\varepsilon) V_{t+1}(a',y',\varepsilon',\alpha,\beta) \right\} \tag{5}$$

$$s.t.$$

$$c + a' = \left(\frac{1+(1-\tau_a)r}{1-\pi_t^d}\right) a + (1-\tau_w-\tau_p)\, wh\epsilon_{t,\varepsilon,\alpha,\beta} + tr$$

$$y' = [(t-1)y + wh\epsilon_{t,\varepsilon,\alpha,\beta}]/t$$

$$a_0 = 0, \quad c \geq 0, \quad a' > -b, \quad 1 \geq h \geq 0$$

Non-retired agents have to set optimally their consumption/savings plans and their labor supply. They enjoy utility from consumption, suffer disutility from labor, and face some uncertain events in the future. In the next period they can still be alive, and with probability $\pi(\varepsilon',\varepsilon)$ they transit from their current persistent shock ε to the value ε'. Their cumulated pension contributions are updated according to their value up to age $t-1$ and the current labor earnings. These agents pay proportional labor and capital income taxes, τ_a and τ_w, respectively. They also pay a proportional tax τ_p on their labor earnings to finance the pension scheme. Finally, they receive a lump-sum transfer tr, they are born with no wealth and with the average shock $\overline{\varepsilon}$, and they are subject to an exogenous borrowing constraint, $b \geq 0$.

4.2 Recursive Stationary Equilibrium

Definition 1 *For given policies* $\{p_1, p_2; \tau_a, G, TR\}$ *a recursive stationary equilibrium is a set of decision rules,* $\{c_t(a,y,\varepsilon,f), a_t'(a,y,\varepsilon,f), h_t(a,y,\varepsilon,f)\}_{t=1}^{T_R-1}$ *and* $\left\{c_t^R(a,y), a_t^{R\prime}(a,y)\right\}_{t=T_R}^{T}$*, value functions,* $\{V_t(a,y,\varepsilon,f)\}_{t=1}^{T_R-1}$

and $\left\{V_t^R(a,y)\right\}_{t=T_R}^{T}$*, prices* $\{r,w\}$*, proportional taxes* $\{\tau_p,\tau_w\}$ *and a set of stationary distributions,* $\{\mu_t(a,y,\varepsilon,f)\}_{t=1}^{T_R-1}$ *and* $\left\{\mu_t^R(a,y)\right\}_{t=T_R}^{T}$*, such that:*

- Given relative prices $\{r,w\}$, proportional taxes $\{\tau_a,\tau_p,\tau_w\}$ and pension benefits $p(y)$, the individual policy functions $\{c_t(a,y,\varepsilon,f), a_t'(a,y,\varepsilon,f), h_t(a,y,\varepsilon,f)\}_{t=1}^{T_R-1}$, $\left\{c_t^R(a,y), a_t^{R\prime}(a,y)\right\}_{t=T_R}^{T}$ solve the household problems (3)-(4), or (3)-(5), and $\{V_t(a,y,\varepsilon,f)\}_{t=1}^{T_R-1}$, $\left\{V_t^R(a,y)\right\}_{t=T_R}^{T}$ are the associated value functions.
- Given relative prices $\{r,w\}$, K/L solves the firm's problem and satisfies (1)-(2).
- The labor market is in equilibrium, and the labor input L corresponds to the total supply of labor efficiency units

$$L=\sum_{t=1}^{T_R-1}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}\times\mathcal{E}\times\mathcal{F}}\epsilon_{t,\varepsilon,f}h_t(a,y,\varepsilon,f)\,d\mu_t(a,y,\varepsilon,f)$$

- The asset market clears

$$(1+n)K=\sum_{t=1}^{T_R-1}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}\times\mathcal{E}\times\mathcal{F}}a_t'(a,y,\varepsilon,f)\,d\mu_t(a,y,\varepsilon,f)+\sum_{t=T_R}^{T}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}}a_t^{R\prime}(a,y)\,d\mu_t^R(a,y)$$

- The goods market clears

$$\begin{aligned}F(K,L) &= C+I+G= \\ &\sum_{t=1}^{T_R-1}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}\times\mathcal{E}\times\mathcal{F}}c_t(a,y,\varepsilon,f)\,d\mu_t(a,y,\varepsilon,f)+\sum_{t=T_R}^{T}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}}c_t^R(a,y)\,d\mu_t^R(a,y)+(\delta+g_n)K+G\end{aligned}$$

- The government's budget is balanced, that is tax revenues from (capital and labor/pension) income taxation are equal to the sum of government purchases G and total transfers TR

$$\begin{aligned}G+TR &= \sum_{t=1}^{T_R-1}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}\times\mathcal{E}\times\mathcal{F}}\tau_a ra\,d\mu_t(a,y,\varepsilon,f)+\sum_{t=T_R}^{T}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}}\tau_a ra\,d\mu_t^R(a,y)+ \\ &\sum_{t=1}^{T_R-1}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}\times\mathcal{E}\times\mathcal{F}}\tau_w w\epsilon_{t,\varepsilon,f}h_t(a,y,\varepsilon,f)\,d\mu_t(a,y,\varepsilon,f)+\sum_{t=T_R}^{T}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}}\tau_w p(y)\,d\mu_t^R(a,y)\end{aligned}$$

- The proportional tax rate of the pension contributions satisfies

$$\tau_p=\frac{\displaystyle\sum_{t=T_R}^{T}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}}p(y)\,d\mu_t^R(a,y)}{\displaystyle\sum_{t=1}^{T_R-1}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}\times\mathcal{E}\times\mathcal{F}}wh_t(a,y,\varepsilon,f)\,\epsilon_{t,\varepsilon,f}\,d\mu_t(a,y,\varepsilon,f)}$$

- The stationary distributions $\left\{\mu_t\left(a,y,\varepsilon,f\right),\mu_t^R\left(a,y\right)\right\}$ satisfy

$$\mu_{t+1}\left(a',y',\varepsilon',f\right)=\int\psi\left(a,y,\varepsilon,f,t,a',y',\varepsilon'\right)d\mu_t\left(a,y,\varepsilon,f\right) \tag{6}$$

$$\mu_{t+1}^R\left(a',y\right)=\int\psi^R\left(a,t,a'\right)d\mu_t^R\left(a,y\right) \tag{7}$$

In equilibrium the measure of agents in each state is time invariant and consistent with individual decisions, as given by the above two equations (6)-(7), where $\psi\left(.\right)$ and $\psi^R\left(.\right)$ are the transition functions.

- The welfare measure W is utilitarian, i.e. it weights agents' utilities by their mass in the steady-state

$$W=\sum_{t=1}^{T_R-1}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}\times\mathcal{E}\times\mathcal{F}}V_t\left(a,y,\varepsilon,f\right)d\mu_t\left(a,y,\varepsilon,f\right)+\sum_{t=T_R}^{T}\Phi_t\int\limits_{\mathcal{A}\times\mathcal{Y}}V_t^R\left(a,y\right)d\mu_t^R\left(a,y\right) \tag{8}$$

5 Calibration and Empirical Analysis

5.1 Parameterization

The economy under study has several independent parameters: $\rho_\varepsilon,\sigma_\eta,\sigma_\alpha^2,\sigma_\beta^2,\sigma_{\alpha\beta},\gamma,\lambda,\delta,d,g_n,\pi_t^d,e_t,T_R,T$. These parameters are divided into three categories. In the first category there are the parameters describing the stochastic processes, $\rho_\varepsilon,\sigma_\eta,\sigma_\alpha^2,\sigma_\beta^2$, and $\sigma_{\alpha\beta}$: these are are going to be assigned several combinations of values, one for each replication. In the second category there are the four parameters that are calibrated in equilibrium: γ,λ,δ,d. Finally, the third category consists of all the demographics/life-cycle parameters: g_n,π_t^d,e_t,T_R,T. These are borrowed from other studies, and are always kept fixed. Agents become economically active at age 20, retire at age $T_R=66$, and they can live up to $T=101$ years. These are fairly conventional and secondary assumptions. The population growth rate g_n, the survival probabilities π_t^d, and the age profile for the efficiency units e_t are all taken from widely accepted studies, which are heavily used in the calibration of similar models.

Stochastic Process Parameters: The parameters driving the labor income uncertainty in the economy are going to be assigned many different values, one set of value for each realization of the draws from their prior distributions.

In the benchmark case, in order to calibrate the remaining parameters, I rely on the estimates obtained by Guvenen (2009). Accordingly, the variance of the persistent shocks is $\sigma_{\eta,SIP}^2=0.015$ for SIP and $\sigma_{\eta,HIP}^2=0.029$ for HIP, the autocorrelations of these shocks are $\rho_{\varepsilon,SIP}=0.988$ for SIP and $\rho_{\varepsilon,HIP}=0.821$ for HIP, and the variances of the fixed effect (level) are $\sigma_{\alpha,SIP}^2=0.058$ for SIP and $\sigma_{\alpha,HIP}^2=0.022$ for HIP. For the HIP process, the variance of the slope fixed effect is $\sigma_\beta^2=0.00038$, and the correlation between fixed effects is $\sigma_{\alpha\beta}=-0.23$.

Preference/Technology Parameters - SIP (HIP): The share of leisure in the utility function λ is chosen to match a value 0.33 for the share of time devoted to work, obtained when $\lambda=0.511$ ($\lambda=0.412$). The concavity of the utility function is pinned down by the IES coefficient, which is set to $\gamma=2.78$ ($\gamma=3.05$) to match a value for the CRRA of 2.0, a common value in the literature. The capital depreciation rate is set to replicate an investment/output ratio of approximately 25.0%. This is achieved with $d = 0.0672$ ($d = 0.0675$). The rate of time preference δ is calibrated to get an equilibrium interest rate of approximately 4% on an annual basis, obtained when $\delta = 0.9705$ ($\delta = 0.9812$), with corresponding capital/output ratios of ≈ 3. Notice how the more intense precautionary saving motive under SIP requires less patient agents to match the same interest rate. With a Cobb-Douglas production function the capital share is captured by the parameter $\phi = 0.36$, which matches the capital share of income. The complete parameterization of the model is reported in Table (2).

[Table 2 about here]

Demographics/Life-cycle Parameters: The US long-run average of the yearly population growth rate is 1.1%, hence $g_n = 0.011$. For the life-cycle aspects of the model, I consider standard profiles for both the survival probabilities, from Bell and Miller (2002), and for the efficiency units e_t, from Hansen (1993).

[Figures 5 and 6 about here]

Figures (5) and (6) plot the former and the latter variables, together with the mass of each cohort, in the bottom panel of Figure (5).

Policy Parameters: Differently from the parameters above, τ_a, G, TR and (p_1, p_2) are policy parameters. They are pinned down by the institutional features of the economy they are meant to represent (its fiscal policies and social security system, in particular). Their values are reported in Table (3).

[Table 3 about here]

The capital income tax rate is 30%, $\tau_a = 0.3$, the public expenditure/output ratio $\left(\frac{G}{Y}\right)$ is approximately 20%, $\frac{G}{Y} = 0.2$, and the public transfers/output ratio $\left(\frac{TR}{Y}\right)$ is assumed to be 5%, $\frac{TR}{Y} = 0.05$. The latter term captures in a parsimonious way the progressivity of the income tax code, together with the redistribution intrinsic in many other insurance and safety net programs, such as Medicare/Medicaid and the food stamps ones, as discussed in Aiyagari and McGrattan (1998). This allows the agents to have the appropriate disposable income. Finally, Section 3 discussed in detail how the values for (p_1, p_2) were chosen.

5.2 The Missing Link: Model Vs. Data

Consider the two economic models (SIP and HIP) in turn. According to the minimal econometric interpretation proposed by Geweke (2010), the models inform us about population moments. Since either economic model has a unique stationary distribution and a unique equilibrium, for a given parameterization they deliver a point for each moment of interest. By assuming priors for the parameters we can obtain the predictive distributions, namely the distributions of the endogenous variables (and their moments) induced by the parameter uncertainty. In this framework there is the need to link the distributions of population moments to the observables. I specify two autocorrelations (ρ_y, ρ_h) as the moments of interest, ρ_y being the autocorrelation of labor earnings and ρ_h being the autocorrelation of hours worked. As for the observables, earnings $(y_{i,t})$ and hours $(h_{i,t})$ are going to be used in my specification.

The choice of this aspect of the analysis is not trivial. On the one hand, we would like to use many different outcomes to evaluate the relative performance of the two models. On the other hand this makes the comparison more complicated to implement, and raises other issues. In principle, since the models are solved in a steady state and have household-level heterogeneity, I could specify empirical models for cross sectional data. However, since the main difference between SIP and HIP pertains to the wage dynamics, it is reasonable to focus the attention on moments dealing with labor earnings and their autocorrelations. In the PSID data I rely on, both annual earnings and annual hours worked are observed for up to 29 years, making them useful sources of information.

A Bayesian Panel VAR: I Specify a Bayesian Panel VAR in labor income $y_{i,t}$ and hours worked $h_{i,t}$, using PSID data on household heads only, for the period 1968-1996.

$$\begin{pmatrix} y_{i,t} \\ h_{i,t} \end{pmatrix} = \begin{bmatrix} \rho_y & c_1 \\ c_2 & \rho_h \end{bmatrix} \begin{pmatrix} y_{i,t-1} \\ h_{i,t-1} \end{pmatrix} + \begin{pmatrix} \eta^y_{i,t} \\ \eta^h_{i,t} \end{pmatrix}$$

$$\begin{pmatrix} \eta^y_{i,t} \\ \eta^h_{i,t} \end{pmatrix} \overset{iid}{\sim} N(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{bmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{bmatrix})$$

All seven parameters, $\rho_y, \rho_h, c_1, c_2, \sigma_{11}, \sigma_{12}, \sigma_{22}$, are estimated with a posterior simulator.[4]

As for the prior distributions for the economic models' parameters, I am assuming uniform priors over the 95% confidence intervals for the SIP and HIP minimum distance estimates on PSID data. These are reported in Table 4.

[Table 4 and Figure 7 about here]

The comparison of the two models is based on the comparison between the (population) moments and the posteriors for the two autocorrelations (ρ_y, ρ_h) obtained from the Bayesian Panel VAR. The Bayes factor for this application is the ratio of the two marginal likelihoods, $P(HIP|y_{i,t}, h_{i,t}, VAR)$ and $P(SIP|y_{i,t}, h_{i,t}, VAR)$ whose expressions are as follows:

[4] Draws of (ρ_y, ρ_h) from the posterior distribution were obtained using a Metropolis within Gibbs posterior simulation algorithm. Improper (flat) priors for all the parameters of the Panel VAR, subject to stationarity conditions, were used. Geweke (2010) carefully explains why flat priors are needed in this framework. The first 10,000 draws are discarded as a burn-in.

$$\frac{P(HIP|y_{i,t},h_{i,t},VAR)}{P(SIP|y_{i,t},h_{i,t},VAR)} \propto \frac{\int\int P(\rho_y,\rho_h|HIP)P(\rho_y,\rho_h|y_{i,t},h_{i,t},VAR^*)d\rho_y d\rho_h}{\int\int P(\rho_y,\rho_h|SIP)P(\rho_y,\rho_h|y_{i,t},h_{i,t},VAR^*)d\rho_y d\rho_h}$$

$$\approx \frac{\frac{1}{N_{HIP}N_{VAR}}\sum_{u=1}^{N_{HIP}}\sum_{v=1}^{N_{VAR}} K\left(\rho_{y,u}^{HIP},\rho_{h,u}^{HIP};\rho_{y,v}^{VAR},\rho_{h,v}^{VAR}\right)}{\frac{1}{N_{SIP}N_{VAR}}\sum_{u=1}^{N_{SIP}}\sum_{v=1}^{N_{VAR}} K\left(\rho_{y,u}^{SIP},\rho_{h,u}^{SIP};\rho_{y,v}^{VAR},\rho_{h,v}^{VAR}\right)}$$

In the formula approximating the double integrals, I assume bivariate normal kernels, and the numbers $N_{SIP}, N_{HIP}, N_{VAR}$ stand for the number of draws for the SIP and HIP models (500 each) and for the Bayesian Panel VAR (10,000).

Figure 7 displays three scatter plots, one for the Bayesian Panel VAR (the left panel), one for the SIP model (the middle panel), and one for the HIP model (the right panel). The first scatter plot displays the draws from the posterior distribution for (ρ_y, ρ_h), while the other two plots display the population autocorrelations related to the SIP and HIP models.

In terms of the results, the most important finding is related to the Bayes Factor. This seems to give more weight to the HIP, but not overwhelmingly so, the marginal likelihood being 3.09 for HIP and 1.98 for SIP, with the Bayes Factor being $3:1$ in favor of the HIP model.

6 Counterfactual Analysis: the Optimal Pension Scheme

In order to characterize the optimal pension scheme it is necessary to take a stand on what constitutes the objective function. To begin with, I will characterize numerically the optimal pension benefits scheme by considering 1) SIP and HIP separately, 2) the average steady-state utility W as the welfare measure, and 3) the pair (p_1, p_2) that achieves the highest welfare as the optimal scheme:

$$\underset{p_1,p_2}{Max} W(p_1,p_2) \equiv \underbrace{\sum_{t=1}^{T_R-1}\Phi_t\int V_t(.)\,d\mu_t(.)}_{Working\text{-}Age} + \underbrace{\sum_{t=T_R}^{T}\Phi_t\int V_t^R(.)\,d\mu_t^R(.)}_{Retirees}$$

For each set of calibrated parameters the equilibrium of the model is computed several times, once for each specific values of the pair (p_1, p_2). Furthermore, notice that also all the counterfactual pension schemes satisfy the balanced budget requirement.

Moreover, notice that in the numerical solution of the model there is no sampling variability (due for example to Monte Carlo integration). The only sources of error are the numerical errors induced by the discretization of the state space, by the interpolation scheme, and by the convergence criteria, which are kept as small as the computational burden makes possible. It follows that, across replications, the computed equilibria vary only because of the different calibrated parameters. Differently, the change in equilibria would partially reflect the simulation error: aggregate quantities would vary randomly, leading to an induced endogenous response of the model and an additional source of error for the welfare measures.

Qualitative analysis: In order to get a better understanding of the mechanisms shaping the optimal pension benefits scheme, it is useful to consider how changes in p_1 and p_2 affect the pension replacement rates.

[**Figures 8 and 9 about here**]

Figure (8) deals with changes in the parameter p_1, while Figure (9) with changes in the parameter p_2. In both Figures, the top panels plot the shape of the PIAR schedule, while the bottom ones the replacement rates. Furthermore, for each panel there are three curves: the green ones represent the shape of the PIAR function at the fitted values for p_1 and p_2, the blue ones the PIAR function after a decrease in either p_1 or p_2 (with the other parameter kept at the benchmark value), and the red ones the PIAR function after an increase in either p_1 or p_2 (with the other parameter kept at the benchmark value).

From Figure (8) it is apparent that increases in p_1 lead to an overall increase in both the PIAR and the replacement rates, with the latter growing more than proportionally. Changes in p_1 have a level effect, and a higher p_1 leads to higher replacement rates for everyone. The redistribution takes place through the contributions, because pensions are financed with a proportional tax rate, while the replacement rate is highly progressive. From an intuitive perspective, the heterogeneity due to the persistent shocks matters most in determining this parameter, as the higher the uncertainty agents face in the labor market, the higher the demand for ex-post redistribution.

From Figure (9) it is apparent that increases in p_2 change the concavity of both the PIAR and the replacement rates, with the latter growing more than proportionally. For a given p_1, changes in p_2 have a direct redistribution effect. Namely, there is a tension between the rich and poor, in terms of their pension contributions. This also explains why I rescaled the PIA by the average contributions $\overline{y}$. With this formulation, an increase in p_2 leads to higher replacement rates below the mean but also to lower replacement rates above it (and vice versa when p_2 decreases). This not only allows for more flexibility in the pension benefits, but also to exacerbate different demands for redistribution, with the below average contributors to the system being able to extract higher benefits from a given pool of contributions from the above average contributors. From an intuitive perspective, the heterogeneity due to the fixed effects matters most in shaping this parameter.

6.1 Optimal Pension System (with only one parameter: p_1)

This computational experiment deals with finding the optimal pension benefits scheme when the optimization is performed as a function of the parameter p_1 only.

Distributions of the optimal p_1: Figure (10) plots the non-parametric kernel densities of the optimal p_1^* obtained for the SIP and HIP models.

[Figures 10 and 11, and Table 5 about here]

The first consideration arising from inspecting the supports and the shape of the distributions is that SIP implies consistently higher pension benefits. A second consideration is that extremely high or low values for the optimal p_1^* are rare, with the $[0.35, 0.47]$ range having most of the mass for HIP, and the $[0.52, 0.62]$ one for SIP.

Table (5) reports several statistics for the optimal p_1^*. The mean value for SIP (HIP) is 0.59 (0.43), the maximum is 0.66 (0.55), while the minimum is 0.49 (0.26).

The optimal values of the parameter p_1^* are not terribly informative. However, the schedules representing the pension replacement rates can be easily computed. Figure (11) displays the replacement rates for both SIP and HIP corresponding to their average p_1^*. The qualitative behavior is similar, with high replacement rates granted to retirees with low contributions, unlike the quantitative findings that differ markedly. As expected, SIP commands a replacement rate which is always higher compared to HIP, irrespective of the pension contributions. In particular, for the average optimal scheme for SIP, the retirees with the lowest contributions (which are approximately 10% of the average ones) have a 101% replacement rate. The corresponding figure for HIP is only 56%. Finally, the optimal (average) replacement rate with HIP is higher than what is typically found in similar studies (i.e., 36% Vs. 30%), but lower than just with SIP (i.e., 47% Vs. 36%).

7 Discussion and Conclusions

This paper explored the use of Bayesian methods to parameterize GE models with heterogeneous agents, which are a useful tool for the ex-ante evaluation of policy changes. More robust methods for their empirical implementation have become feasible and in this paper I have developed an empirical framework robust to both parameter uncertainty and uncertainty over labor market risk. It was found that the optimal replacement rate of the social security system depends heavily on the nature of labor income risk. Preliminary results seem to give more weight to the HIP, but not overwhelmingly so (with a Bayes Factor of $\approx 3:1$). Finally, the optimal replacement rate with HIP was found to be higher than what is typically found in similar studies (i.e., 36% Vs. 30%), but lower than just with SIP (i.e., 47% Vs. 36%).

The next step of the analysis is to perform the optimization with respect to both parameters characterizing the PIA function. Given that the parameter p_2 governs the demand of redistribution between individuals with high vs. low contributions to the pension system, it is hard to predict what the optimal scheme will look like.

A learning mechanism between SIP and HIP for the households would represent a desirable feature of the model. With the current specification, the individuals know which process is driving their wage shocks, unlike the government. However, the additional computational complexity might make it unfeasible. Moreover, notice that the current framework is still coherent, as any attempt of estimating the DGP with state of the art econometric models but samples of limited size subject to measurement error by public authorities would lead to non-overwhelming evidence. Other issues are whether to incorporate more moments in the empirical validation of the model, and whether to attempt a full-blown Bayesian analysis (which is unlikely to be feasible). Finally, the role of public debt could be included in the framework, as Aiyagari and McGrattan (1998) have shown that the optimal quantity of public debt is positive and large, even in the absence of a public pension system.

It goes without saying that a similar framework can be used to evaluate other policies, such as the optimal capital income taxation, or the tax rebates aimed at stimulating the aggregate activity, such as the recent Fiscal Stimuli.

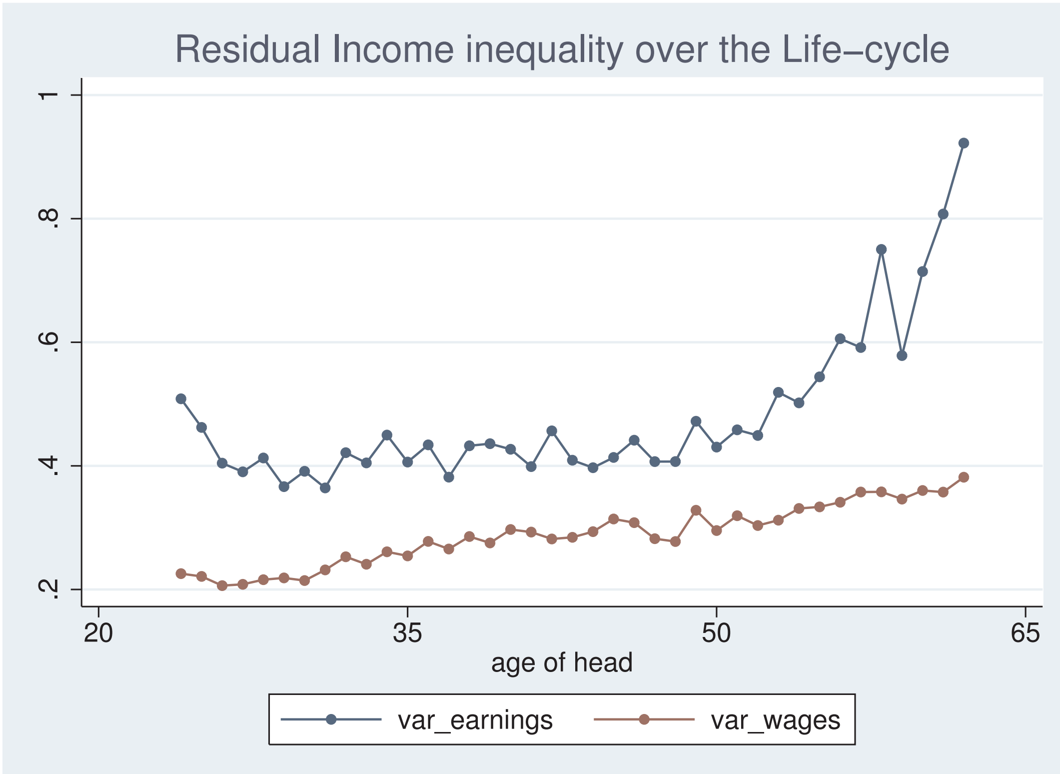

Figure 1: (Residual) Variances of Log-Wages and Log-Earnings, PSID 1968-1996.

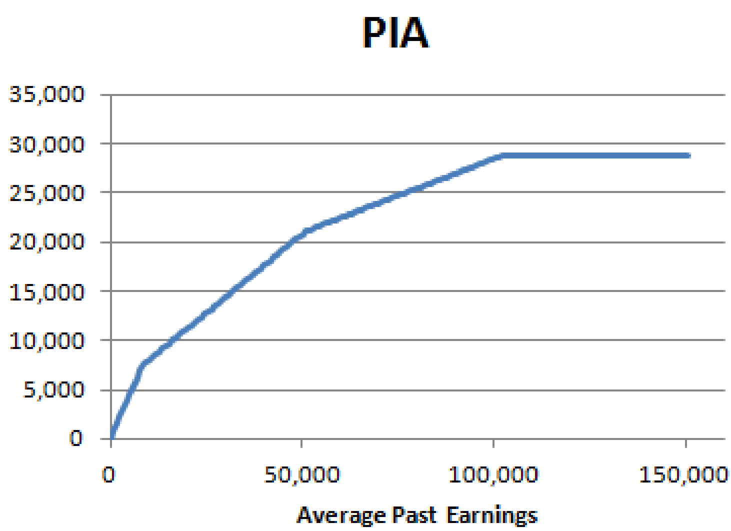

Figure 2: Pension Benefits in levels ($2008)

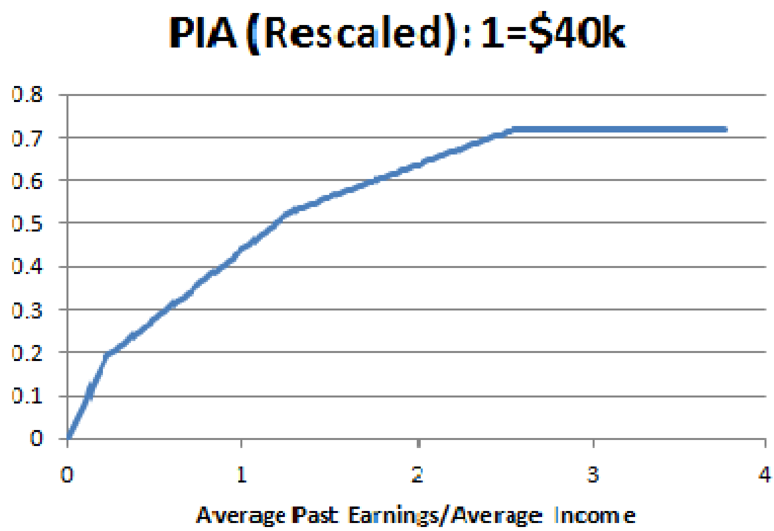

Figure 3: Pension Benefits rescaled by Avg income

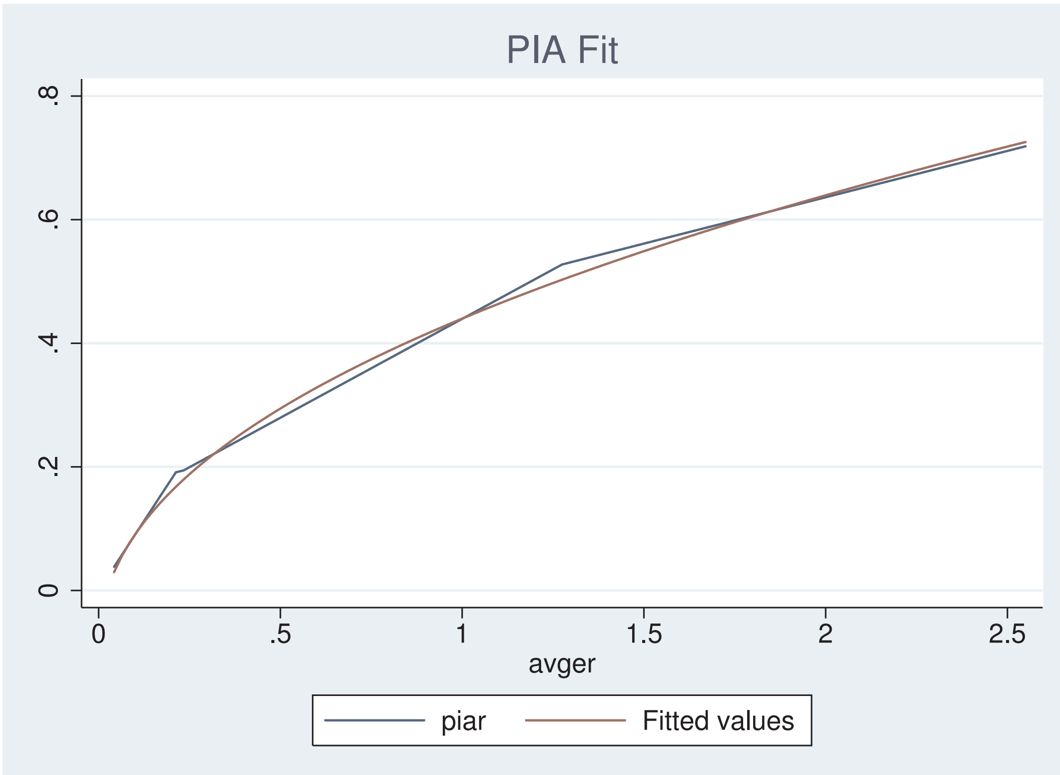

Figure 4: Actual PIA Vs. Fitted one.

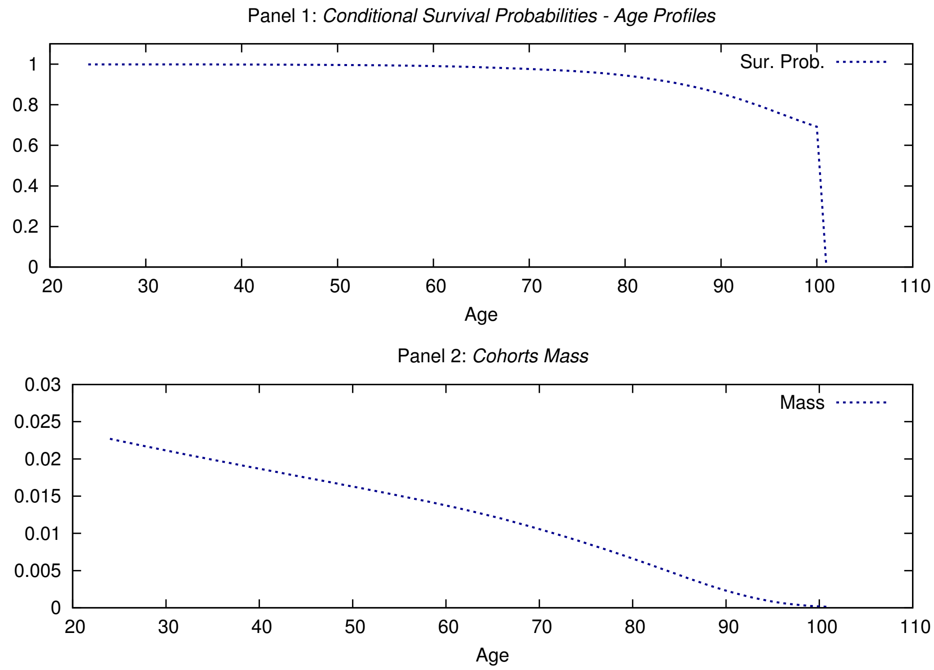

Figure 5: Life-Cycle Profiles: Survival Probabilities (Bell-Miller 2002) and Cohorts Mass.

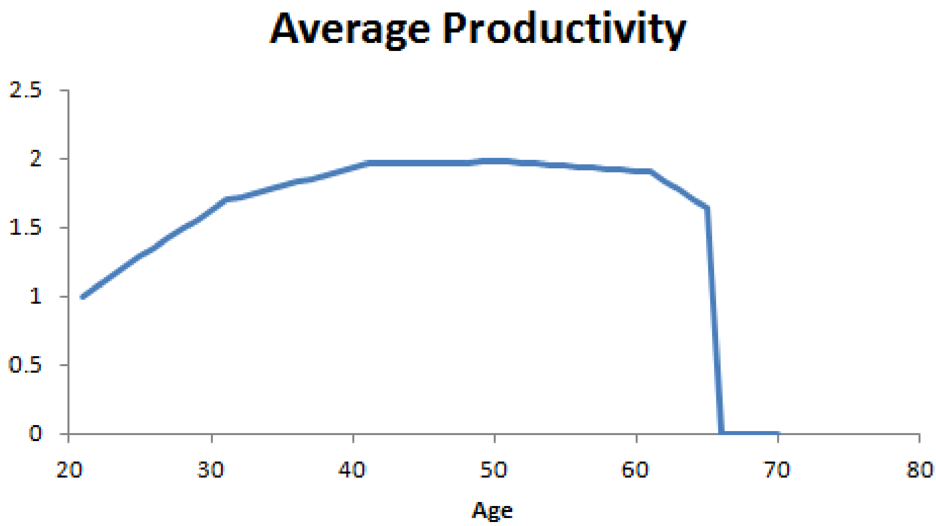

Figure 6: Average Efficiency units over the Life-cycle (Hansen 1993).

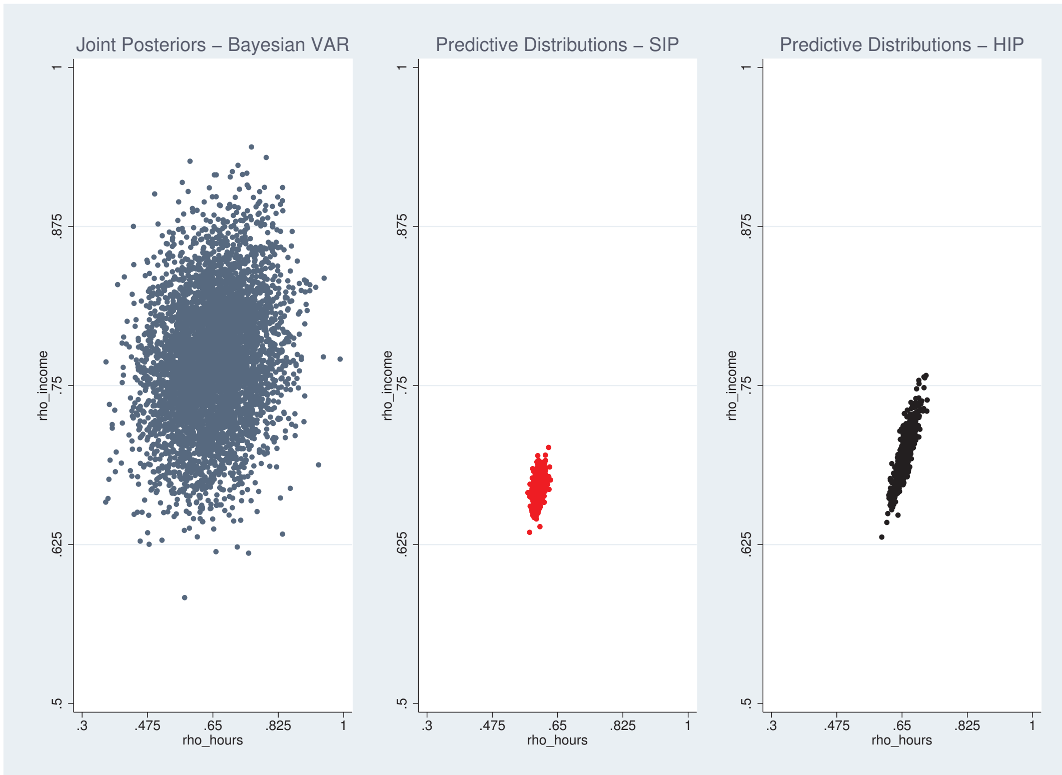

Figure 7: Linking the Models to the Data: Autocorrelations of Labor Income and Hours. *Log Marginal Likelihoods*: 3.09 (HIP) Vs. 1.98 (SIP). *Bayes Factor* 3:1.

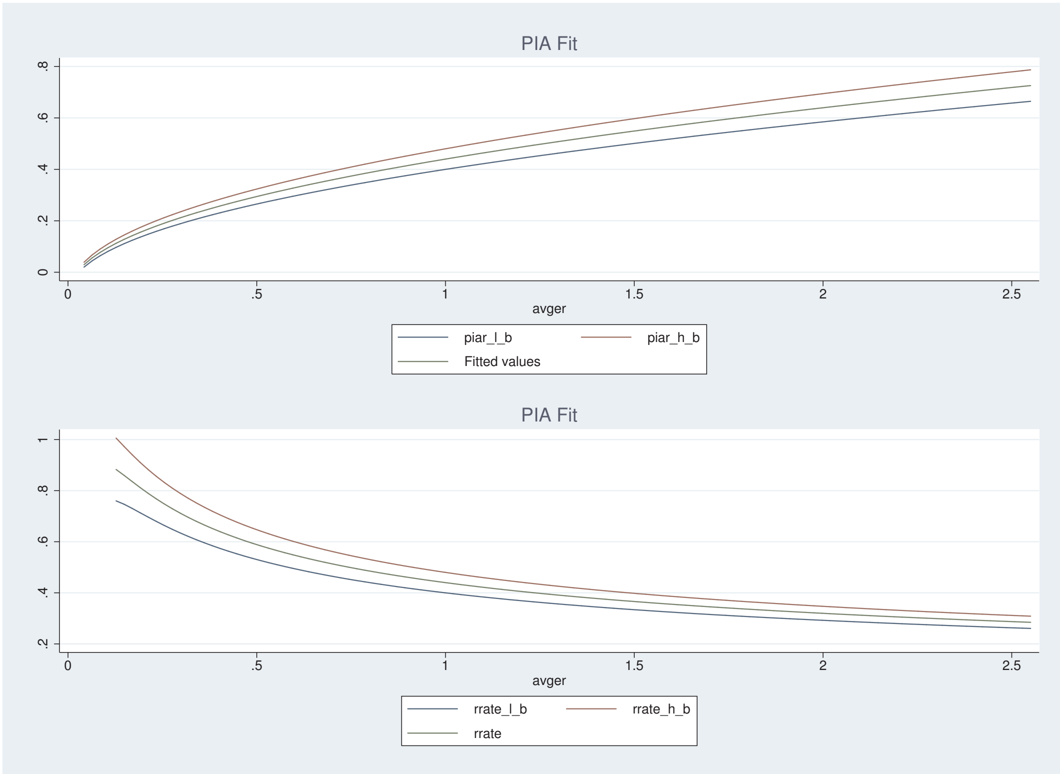

Figure 8: Qualitative behavior, p_1. Top panel: rescaled benefits; Bottom panel: replacement rates.

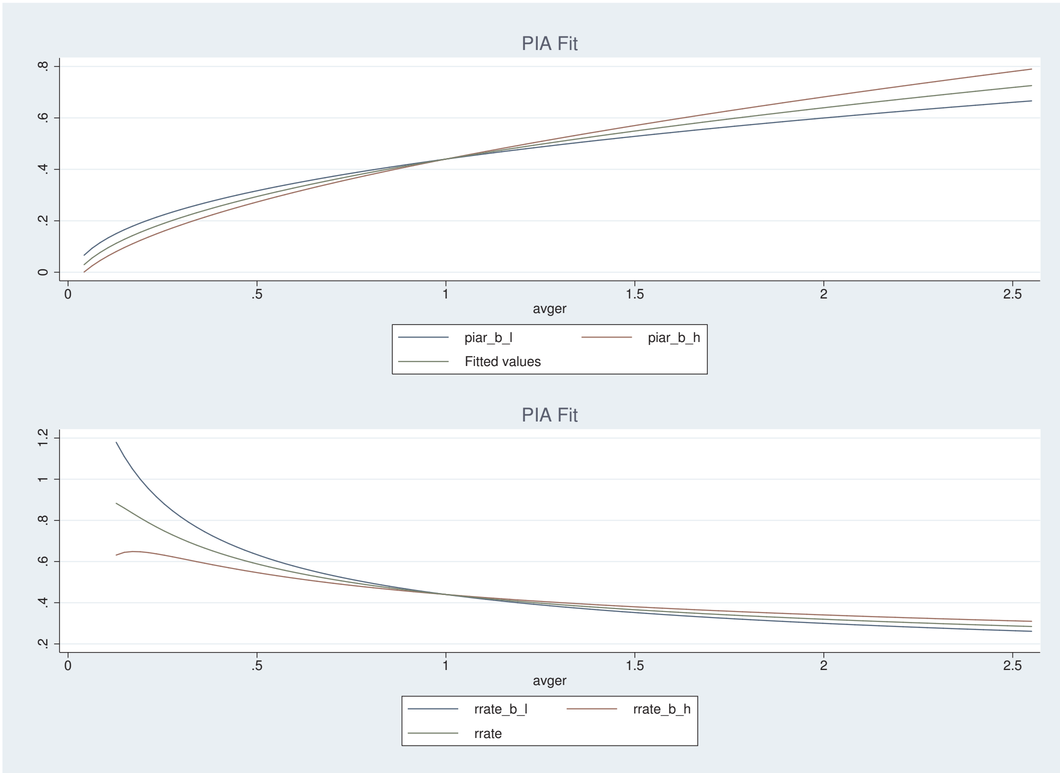

Figure 9: Qualitative behavior, p_2. Top panel: rescaled benefits; Bottom panel: replacement rates.

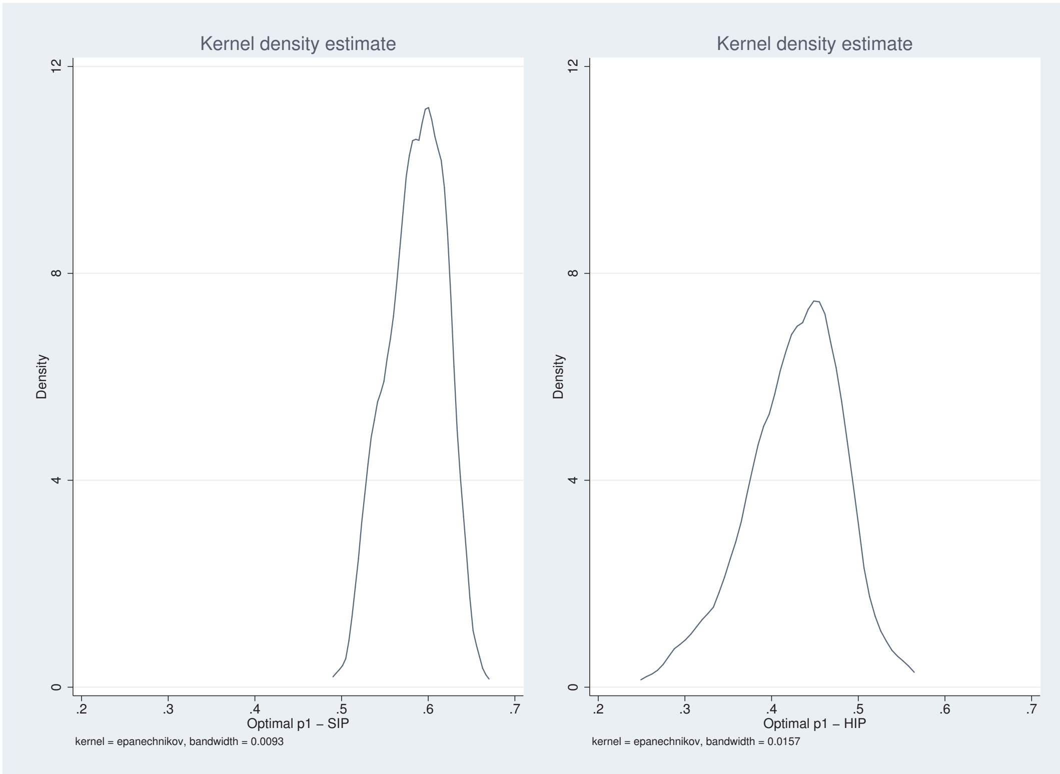

Figure 10: Results - Optimal Social Security (p_1, 500 replications): SIP (left panel) Vs. HIP (right panel).

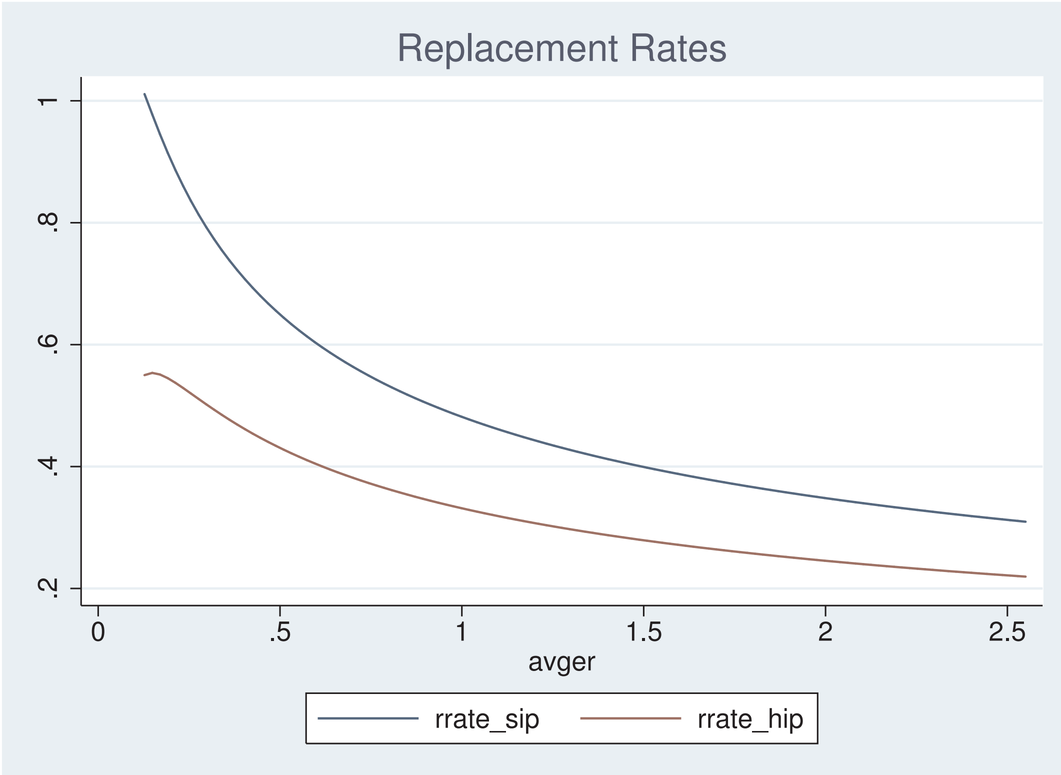

Figure 11: Optimal replacement rates: SIP Vs. HIP evaluated at the average p_1^*.

Parameter		*SIP*	*HIP*
σ_η^2	*Variance of the persistent shocks*	0.015	0.029
ρ_ε	*Autocorrelation of the persistent shocks*	0.988	0.821
σ_α^2	*Variance of the FE (level)*	0.058	0.022
σ_β^2	*Variance of the FE (slope)*	$-$	0.00038
$\sigma_{\alpha\beta}$	*Correlation between FE's*	$-$	-0.23

Table 1: Parameterization, Labor Market Risk - Wage Processes Minimum Distance Estimates on PSID data. Source: Guvenen (2009).

Parameter	*SIP*	*HIP*	*Target*
Model Period	*Yearly*	*Yearly*	
In Equilibrium:			
δ - *Discount factor*	0.9705	0.9812	*Interest rate* = 4%
λ - *leisure share*	0.511	0.412	Work time = 0.33
γ - $1/IES$	2.78	3.05	*CRRA* = 2
d - *Capital deprec.*	0.0672	0.0675	*I/Y ratio* = 25%
Fixed:			
ϕ - *Capital share*	0.36	0.36	*Capital share of Y*
n - *pop. growth*	0.011	0.011	$1970 - 2010$ *average*
e_t - *productivity*	*See fig.*	*See fig.*	*Hansen 1993*
π_t^d - *death prob.*	*See fig.*	*See fig.*	*Bell-Miller 2002*

Table 2: Benchmark Calibration, under current PIA in the US.

Parameter	*Value*	*Policy*
τ_a	0.30	*Capital income tax*
$(\overline{G/Y})$	0.20	*Government expenditure (20% of Steady State Output)*
$(\overline{TR/Y})$	0.05	*Lump-sum public transfers (5% of Steady State Output)*
p_1 - *PIA level*	0.538	*Current PIA function*
p_2 - *PIA concavity*	0.455	*Current PIA function*

Table 3: Policy Parameters, common to both SIP and HIP Model.

	Parameter	*SIP*	*HIP*
σ^2_η	*Variance of the persistent shocks*	$U[0.001, 0.029]$	$U[0.013, 0.045]$
ρ_ε	*Autocorrelation of the persistent shocks*	$U[0.94, 1.0]$	$U[0.761, 0.881]$
σ^2_α	*Variance of the FE (level)*	$U[0.036, 0.07]$	$U[0, 0.17]$
σ^2_β	*Variance of the FE (slope)*	$-$	$U[0.00022, 0.00054]$
$\sigma_{\alpha\beta}$	*Correlation between FE's*	$-$	$U[-0.66, 0.20]$

Table 4: Prior Distributions for the two Models' Stochastic Income Processes Parameters: 95% confidence intervals around the classical estimates.

	Min	*Max*	*Mean*	*Med*	*S.d.*
HIP	.26	.55	.43	.43	.055
SIP	.49	.66	.58	.59	.032

Table 5: Optimal Social Security: Statistics of p_1 over 500 replications.

References

Aiyagari, R. (1994). "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics,* Vol. 109, 659-684.

Aiyagari, R., and McGrattan, E. (1998). "The optimum quantity of debt," *Journal of Monetary Economics,* Vol. 42, 447-469.

Bell, F., and Miller, M. (2002). "Life Tables for the United States Social Security Area 1900–2100." Office of the Chief Actuary, Social Security Administration, Actuarial Study 116.

Canova, F. (1994). "Statistical Inference in Calibrated Models," *Journal of Applied Econometrics*, Vol. 9, S123-S144.

Canova, F. (1995). "Sensitivity Analysis and Model Evaluation in Simulated Dynamic General Equilibrium Economies," *International Economic Review*, Vol. 36 , 477-501.

Conesa, J., Kitao, S. and Krueger, D. (2009). "Taxing Capital? Not a Bad Idea after All!," *American Economic Review*, Vol. 99, 25-48.

Cozzi, M. (2012). "Optimal Unemployment Insurance in GE: a Robust Calibration Approach," *Economics Letters*, Vol. 117, 28-31.

French, E. (2005). "The Effects of Health, Wealth and Wages on Labor Supply and Retirement Behavior," *Review of Economic Studies*, Vol. 72, 395-427.

Fuster, L., Imrohoroglu, A., and Imrohoroglu, S. (2003). "A Welfare Analysis of Social Security in a Dynastic Framework," *International Economic Review*, Vol. 44, 1247-1274.

Fuster, L., Imrohoroglu, A., and Imrohoroglu, S. (2007). "Elimination of Social Security in a Dynastic Framework," *Review of Economic Studies*, Vol. 74, 113-145.

Geweke, J. (2010). *Complete and Incomplete Econometric Models*. Princeton University Press, Princeton (NJ).

Gouveia, M., and Strauss, R. (1994). "Effective Federal Individual Income Tax Functions: An Exploratory Empirical Analysis," *National Tax Journal*, Vol. 47, 317-39

Guvenen, M.F. (2009). "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, Vol. 12, 58-79.

Guvenen, F. (2007). "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?," *American Economic Review*, Vol. 97, 687-712.

Hansen, G. (1993). "The Cyclical and Secular Behaviour of the Labour Input: Comparing Efficiency Units and Hours Worked." *Journal of Applied Econometrics*, Vol. 8, 71-80.

Heathcote, J., Perri, F., and Violante, G. (2010). "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006," *Review of Economic Dynamics*, Vol. 13, 15-51.

Heathcote, J, Storesletten, K., and Violante, G. (2009). "Quantitative Macroeconomics with Heterogeneous Households," *Annual Review of Economics*, Vol. 1, 319-354.

Heathcote, J, Storesletten, K., and Violante, G. (2010). "The Macroeconomic Implications of Rising Wage Inequality in the United States," *Journal of Political Economy*, Vol. 118, 681-722.

Huggett, M. (1996). "Wealth distribution in life-cycle economies," *Journal of Monetary Economics*, Vol. 38, 469-494.

Huggett, M., Ventura, G., and Yaron, A. (2011). "Sources of Lifetime Inequality," *American Economic Review*, Vol. 101, 2923-2254.

Imrohoroglu, A., Imrohoroglu, S., and Joines, D. (1998). "The Effect of Tax-Favored Retirement Accounts on Capital Accumulation," *American Economic Review*, Vol. 88, 749-768.

Imrohoroglu, A., Imrohoroglu, S., and Joines, D. (2003). "Time-Inconsistent Preferences and Social Security," *Quarterly Journal of Economics*, Vol. 118, 745-784.

Imrohoroglu, S., and Kitao, S. (2012). "Social Security Reforms: Benefit Claiming, Labor Force Participation and Long-run Sustainability," *American Economic Journal: Macroeconomics*, Vol. .

Kopecky, K., and Suen, R. (2010). "Finite state Markov-chain approximations to highly persistent processes," *Review of Economic Dynamics*, Vol. 13, 701-714.

Rios-Rull, J.V. (1999). "Computation of Equilibria in Heterogenous Agent Models," in *Computational Methods for the Study of Dynamic Economies: An Introduction*, (Ramon Marimon and Andrew Scott, Eds.), Oxford University Press.

Storesletten, K., Telmer, C. and Yaron, A. (2004a). "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics*, Vol. 51, pp. 609-633.

Storesletten, K., Telmer, C. and Yaron, A. (2004b). "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, Vol. 112, 695-717.

Appendix A - Computation

- All codes solving the model economies were written in the FORTRAN 2003 language, relying on the Intel Fortran Compiler, build 13 (without the IMSL library). They were compiled selecting the O3 option (maximize speed), and without automatic parallelization. They were executed on a 64-bit PC platform, running Windows 7 Professional Edition, with an Intel $i7-2600k$ Quad Core processor clocked at 4.6 Ghz. They are currently being migrated to a server with a 64-bit Linux platform, running CentOS 6.4, with a dual Intel Xeon $E5-2687W$ octo-core 3.1 Ghz processor, and with eight Intel Xeon Phi $3120P$ co-processors, parallelizing the different calibrations on the co-processors with the OpenMP directives.

- The 500 replications currently take up to 600 hours to complete. Notice that $500 \times J$ equilibria have to be computed, with J representing the total grid size for the variables representing the PIA scheme. Typically from 10 to 14 iterations on the endogenous variables are needed to find each equilibrium.

- In the actual solution of the models I need to discretize the continuous state variables a, ε, y, f (note: f stands for the fixed effect components, i.e. α for SIP and α, β for HIP.) As for a I rely on an unevenly spaced grid, with the distance between two consecutive points increasing geometrically. This is done to allow for a high precision of the policy rules at low values of a, where the change in curvature is more pronounced. I use 201 points, as increasing the number of points does not affect the results considerably, the lowest value being the borrowing constraint and the highest one being a value $a_{\max} > \overline{a}$ high enough for the saving functions to cut the 45 degree line ($a_{\max} = 150$). ε is discretized with the Rouwenhorst method, using an 11-state Markov chain. This method has several desirable properties, especially when working with highly persistent processes, as discussed by Kopecky and Suen (2010). The grids over the fixed effects are obtained by discretizing a normal variate, for the SIP, or a correlated bivariate normal, for the HIP. I use 6 points for α and 6 points for β. As for the pension contributions, captured by the average labor income, I use 10 points.

- The model is solved with a backward recursion on the Bellman equations. I start from the terminal value $V_{T+1}^{R} = 0$, and at each age, for every point in the state space, I solve the constrained maximization problem. I retrieve the policy functions, $a_t'(a, y, \varepsilon, f), h_t(a, y, \varepsilon, f)$ and $a_t^{R\prime}(a, y)$, and I compute the vector of parameters Ω representing the Schumaker spline approximations of the value functions. Notice that I do not restrict the agents' asset holdings and the pensions contributions to belong to discrete sets. As for the approximation method, I rely on the quadratic spline approximations for the future value functions, when evaluated at the chosen saving level.

- I do not use a sample of individuals to approximate the stationary distributions. The distributions are computed relying on their definitions (6)-(7). I rely on these recursions and compute numerically the transitions functions.

- The asset market is in equilibrium when the current guess for the interest rate r_0 achieves a capital excess demand which is less than 0.1% of the market size. In turn, this implies that the excess demand in the final good market is always less than 0.1% of the market size.

- The welfare measures W_{SIP}, W_{HIP} are just the numerical integral of the value functions, integrated with respect to the steady state distributions.

Appendix B - Algorithms

This algorithm represents the computational procedure used to solve the two OLG models with parameter uncertainty.

1. Calibrate both models under the current US PIA and compute the aggregate ex-ante welfares W_{SIP} and W_{HIP}.
2. Draw n combinations of parameters from their prior distributions and store them.
3. Set the indicator $i = 1$.
4. Parametrize the model's wage process by reading the $i - th$ vector of simulated parameters.
5. Generate a discrete grid over the PIA variables $[p_{1,\min}, ..., p_{1,\max}] \times [p_{2,\min}, ..., p_{2,\max}]$.
6. Set the counterfactual PIA function $(p_1, p_2) = (p_{1,j_1}, p_{2,j_2})$ and begin the model solution.
7. Generate a discrete grid over the asset space $[0, ..., a_{\max}]$.
8. Generate a discrete grid over the income shocks with the Rouwenhorst method $[\varepsilon_{\min}, ..., \varepsilon_{\max}]$.
9. Generate a discrete grid over the fixed effects: $[\alpha_{\min}, ..., \alpha_{\max}]$ (SIP), or $[\alpha_{\min}, ..., \alpha_{\max}] \times [\beta_{\min}, ..., \beta_{\max}]$ (HIP).
10. Generate a discrete grid over the pension contributions $[y_{\min}, ..., y_{\max}]$.
11. Start the loop for the benchmark economy.
12. Guess the labor supply L_0.
13. Guess the interest rate r_0.
14. Get the capital demand K_0 and wages w_0.
15. Guess the aggregate pension benefits P_0.
16. Get the equilibrium taxes τ_p, τ_w.
17. Get the saving functions $a_t' \left(a, y, \varepsilon, f; p_{1,j_1}, p_{2,j_2}\right)_i, a_t^{R\prime} \left(a, y; p_{1,j_1}, p_{2,j_2}\right)_i$, the labor supply function $h(a, y, \varepsilon, f; p_{1,j_1}, p_{2,j_2})_i$ and the value functions $V_t \left(a, y, \varepsilon, f; p_{1,j_1}, p_{2,j_2}\right)_i, V_t^R \left(a, y; p_{1,j_1}, p_{2,j_2}\right)_i$.
18. Get the stationary distributions $\mu_t \left(a, y, \varepsilon, f; p_{1,j_1}, p_{2,j_2}\right)_i, \mu_t^R \left(a, y; p_{1,j_1}, p_{2,j_2}\right)_i$.
19. Get the aggregate capital supply and check the asset market clearing; Get r_1.
20. Update $r_0' = \varpi_r r_0 + (1 - \varpi_r) r_1$ (with ϖ_r an arbitrary weight).
21. Get the aggregate labor supply and check the labor market clearing; Get L_1.
22. Update $L_0' = \varpi_L L_0 + (1 - \varpi_L) L_1$ (with ϖ_L an arbitrary weight).

23. Get the aggregate pension benefits; Get P_1.

24. Update $P_0' = \varpi_P P_0 + (1 - \varpi_P) P_1$ (with ϖ_P an arbitrary weight).

25. Iterate until asset and labor market clearing, and until the government budget constraint is satisfied.

26. Compute the average steady-state utility $W(p_{1,j_1}, p_{2,j_2})_i = \sum_{t=1}^{T_R-1} \Phi_t \int_{\mathcal{A}\times\mathcal{Y}\times\mathcal{E}\times\mathcal{F}} V_t(.; p_{1,j_1}, p_{2,j_2})_i \, d\mu_t(.; p_{1,j_1}, p_{2,j_2})_i +$
$\sum_{t=T_R}^{T} \Phi_t \int_{\mathcal{A}\times\mathcal{Y}} V_t^R(.; p_{1,j_1}, p_{2,j_2})_i \, d\mu_t^R(.; p_{1,j_1}, p_{2,j_2})_i$.

27. Consider the next point in the PIA grid, until $(p_{1,j_1}, p_{2,j2}) = (p_{1,_{\max}}, p_{2,_{\max}})$.

28. Compute the optimal PIA scheme $(p_1^*, p_2^*)_i$ as the steady-state utility maximizer, i.e. $(p_1^*, p_2^*)_i \equiv \underset{p_{1,j_1}, p_{2,j_2}}{Arg\max} \left\{W(p_{1,j_1}, p_{2,j_2})_i\right\}$

29. Save the output, set $i = i + 1$ and repeat steps $4 - 28$ for all the n combinations of simulated parameters.

This algorithm represents the computational procedure used to solve the Monte Carlo experiments under uncertainty about the two OLG models.

1. Draw $2n$ combinations of parameters (n for SIP and n for HIP) from their prior distributions and store them.

2. Draw SIP or HIP from a Binomial distribution, with probability proportional to their Bayes factor.

3. Repeat steps $4 - 27$ above and compute the average steady-state utility $W(p_{1,j_1}, p_{2,j_2})_i$

4. Compute the optimal PIA scheme $(p_1^*, p_2^*)_i$ as the steady-state utility maximizer, i.e. $(p_1^*, p_2^*)_i \equiv \underset{p_{1,j_1}, p_{2,j_2}}{Arg\max} \left\{W(p_{1,j_1}, p_{2,j_2})_i\right\}$