

Optimal Pension Systems with Simple Instruments[†]

By MIKHAIL GOLOSOV, ALI SHOURIDEH, MAXIM TROSHKIN, AND ALEH TSYVINSKI*

We analyze optimal pension systems relying on simple policy instruments. The optimality in this context means the highest welfare that can be achieved with a restricted set of tax instruments. While there is an arbitrary number of ways to restrict a tax system, we are guided by the form of the current US Social Security.

The US social insurance system, Social Security, includes several programs or types of insurance it provides on a mandatory basis. We focus on the largest component—the retirement benefit program. The policy instruments we consider are the optimized retirement benefit functions. Huggett and Parra (2010), in the context of Social Security, and Conesa, Kitao, and Krueger (2009), in the context of dynamic taxation, follow a similar approach.

The main difference between our work and Huggett and Parra (2010) is that we are interested in a more sophisticated underlying setup of the individual decision problem which admits endogenous decisions of how much to work as well as when to retire. Specifically, we consider a model with both intensive and extensive margins of labor supply that are meaningfully active as in Shourideh and Troshkin (2012).

To achieve the full optimum in such an environment, under realistic assumptions, policy-relevant implementations involve nonlinear history-dependent income taxes and retirement benefits that change with the actual retirement

age. Instead, we take the existing US retirement benefit function and compute its optimal version by altering the parameters of the current benefit function. That is, we determine the optimum with restricted instruments.

I. Environment

The environment closely follows Shourideh and Troshkin (2012). Time is continuous from $t = 0$ to $t = 1$. A continuum of individuals are born at $t = 0$ and live until $t = 1$. At every point in life, each individual chooses whether to work or not, and if so, how much. Individuals consume a single consumption good and leisure.

Individuals differ in two respects. First, individual workers are heterogeneous in their productivities. Individual productivity changes over life cycle and follows an idiosyncratic hump-shaped productivity profile. When the individual works l hours at age t , her income is $\varphi(t, \theta)l$. Second, individuals face heterogeneous fixed costs of work $\eta(\theta)$ whenever they work non-zero hours.

Specifically, at $t = 0$, each individual draws a type, θ , from a distribution of types, $F(\theta)$, where $F'(\theta) = f(\theta) > 0$ for all θ . Individual's preferences are then represented by $U(\theta)$ given by

$$(1) \quad \int_0^1 e^{-\rho t} \left[u(c(t, \theta)) - v\left(\frac{y(t, \theta)}{\varphi(t, \theta)}\right) - \eta(\theta) \mathbf{1}[y(t, \theta) > 0] \right] dt$$

over the set of allocations $\{c(t, \theta), y(t, \theta)\}_{t \in [0, 1]}$ of consumption and income for each θ . Here, $\mathbf{1}[y(t, \theta) > 0]$ is an indicator function of positive output. The function $u(\cdot)$ is strictly concave, increasing, and satisfies standard Inada conditions; $v(\cdot)$ is a strictly convex function with $v'(0) = 0$. These preferences exhibit fixed costs of working. Combined with a hump-shaped productivity profile, this makes it optimal

*Goloso: Department of Economics, Princeton University, 111 Fisher Hall, Princeton, NJ 08544 (e-mail: goloso@princeton.edu); Shourideh: Wharton School, University of Pennsylvania, 2335 Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104 (e-mail: shouride@wharton.upenn.edu); Troshkin: Department of Economics, Cornell University, 446 Uris Hall, Ithaca, NY 14853 (e-mail: troshkin@cornell.edu); Tsyvinski: Department of Economics, Yale University, Box 208268, New Haven, CT 06520 (e-mail: a.tsyvinski@yale.edu). We thank Florian Scheuer for a discussion and 2013 AEA session participants for comments.

[†]To view additional materials, and author disclosure statement(s), visit the article page at <http://dx.doi.org/10.1257/aer.103.3.502>.

for a worker to choose to retire at some age $R(\theta) \in [0, 1]$, i.e., discontinuously choose $y(t, \theta) = 0$ for all $t \geq R(\theta)$. The heterogeneity implies that retirement ages differ among workers. In other words, this environment features both active intensive and extensive labor margins in the form of decisions about how much to work and when to retire.

The government maximizes a social welfare function given by

$$(2) \quad \int_{\underline{\theta}}^{\bar{\theta}} U(\theta) dG(\theta),$$

where $U(\theta)$ is the lifetime utility of a household of type θ given by (1). The function $G(\theta)$ is a cumulative density function, i.e., $G(\underline{\theta}) = 0$, $G(\bar{\theta}) = 1$, and $G'(\theta) = g(\theta) \geq 0$, and $G(\theta)$ is differentiable over interval $[\underline{\theta}, \bar{\theta}]$. A redistributive motive for the government implies $G(\theta) \geq F(\theta)$ for all $\theta \in [\underline{\theta}, \bar{\theta}]$. We explore various redistributive motives with the baseline case of a utilitarian planner, i.e., $F(\theta) = G(\theta)$.

We consider a planner that chooses allocations $\{c(t, \theta)\}_{t \in [0, 1], \theta \in [\underline{\theta}, \bar{\theta}]}$, $\{R(\theta)\}_{\theta \in [\underline{\theta}, \bar{\theta}]}$, and $\{y(t, \theta)\}_{t \in [0, R(\theta)], \theta \in [\underline{\theta}, \bar{\theta}]}$ and an optimal pension system (restricted to simple instruments made precise below) to maximize social welfare (2), subject to individual budgets and individual optimality in the presence of the pension system. Before turning to the details of the individual optimality, we describe next a stylized version of the US Social Security retirement benefit function, that guides the set of simple instruments we analyze.

II. Stylized Social Security

The US Social Security's Old-Age, Survivors, and Disability Insurance (OASDI) program is quite complex. We focus on the largest component of it, the retirement benefit program, and provide a stylized description of the OASDI retirement benefit function. Then, based on this description, we construct a parsimonious benefit function that we implement in the model.

The OASDI's retirement benefit function is referred to as the Primary Insurance Amount (PIA). Consider a worker with a sequence of incomes over life cycle given by $\{y_t\}_{t \in [S, R]}$, where S is his age of earliest employment and R^{SS} is

his claimed retirement benefits age.¹ Denote the Social Security early retirement age by R^E and the normal retirement age for the worker's year of birth by R^N .²

Compute PIA as a function of one argument—the average indexed monthly earnings (AIME). AIME is 1/12 of the mean of the 35 highest y_t in real terms. Obtain the real terms by inflating by the national average wage index (AWI) up to the year when the worker is $R^E - 2$ years old. If the income in period t , y_t , is greater than OASDI benefits wage base, replace y_t with the wage base in that year.³ If a worker did not work a year set $y_t = 0$. If a worker has less than 35 years of wages, $S - R^{SS} + 1 < 35$, add zeros to get to 35.

PIA is then a piecewise-linear function of AIME with three segments with marginal rates (i.e., slopes) of 0.9, 0.32, and 0.15 (the marginal benefit rate is zero above the benefit base). The segment cutoffs are referred to as PIA bend points.⁴ That is, a worker gets 0.9 of AIME up to the first bend point, plus 0.32 up to the second bend point, plus 15 percent up to the wage base, and nothing more if the AIME is higher still.

The monthly PIA retirement benefits are adjusted as follows. If $R^E \leq R^{SS} < R^N$: the PIA is reduced for each month below R^N , first 36 months by 5/900, then by 5/1,200 (depending on R^N this gives at R^E 70–80 percent of PIA). If $R^{SS} > R^N$: the PIA is increased by 8 percent for each year above R^N up to 132 percent of PIA or up to age 70, whichever comes first. Depending on the year of birth, the cost-of-living adjustments may also be added.

We model PIA benefit function parsimoniously as a function of one argument, AIME, with 6 parameters: a piecewise-linear function starting from zero with two bend points (2 parameters), three slopes (3 parameters), and the benefit base (1 parameter), i.e., the upper bound after which the marginal benefit is zero. Following the current system, we restrict the slopes to be less

¹ Note that in relation to the model, this age, R^{SS} , may be different from the age when the worker decides to actually stop working, R .

² The ages R^N depend on the year of birth. R^N currently ranges 65–67; R^E is equal to 62.

³ For 2013 the wage base is \$113,700. The wage base typically increases every year by anywhere between 0 percent and 6 percent.

⁴ PIA bend points change over time. For 2013 they are \$791 and \$4,768.

than one and non-increasing. Later, we remove the restriction of a zero intercept and study the implications of choosing it optimally. To account for not necessarily actuarially-fair benefit adjustments, we assume that every worker takes advantage of all possible adjustments up to 132 percent as explained above. We assume benefits do not depend on R . In the model, AIME is the working life average income.

To summarize, we model a set of simple instruments for a pension system—a stylized version of the Social Security retirement benefit. Admittedly, the existing system is much more complicated.⁵ However, we believe this description captures the key conceptual elements of the system.

III. Reforms via Simple Policy Instruments

Consider now one of the individuals in our environment who makes decisions about how much to work and when to retire in the presence of the benefit system described above. Upon retiring, the individual is entitled to the present value of lifetime pension benefits as a function of her income profile over working life, $b(Y(\{y(t)\}_{t=0}^R))$. We think of $Y(\cdot)$ as a measure of lifetime income modeling AIME. The function $b(\cdot)$ then models PIA. The individual also faces an effective income tax function $T(\cdot)$. The individual problem is then

$$\begin{aligned} \max_{c(t), y(t), R, a(t)} & \int_0^1 e^{-\rho t} u(c(t)) dt \\ & - \int_0^R e^{-\rho t} \left[v\left(\frac{y(t)}{\varphi(t)}\right) + \eta \right] dt \end{aligned}$$

subject to

$$\begin{aligned} c(t) + \dot{a}(t) &= (y(t) - T(y(t))) \mathbf{1}[t \leq R] \\ &+ \mathbf{1}[t > R^{SS}] \frac{rb(Y(\{y(t)\}_{t=0}^R))}{e^{-rR^{SS}} - e^{-r}} + ra(t), \end{aligned}$$

where $a(t)$ is the level of asset holdings by the individual at date t and $\frac{rb(Y(\{y(t)\}_{t=0}^R))}{e^{-rR^{SS}} - e^{-r}}$ is the

⁵ We do not consider nuanced spousal benefit interactions or survivor benefits and many other features.

level of retirement benefit that the individual receives at every point in time from age R^{SS} to 1, which generates the present value of lifetime benefits equal to $b(Y(\{y(t)\}_{t=0}^R))$.

We rewrite the date-by-date budget constraints above as the following present value budget constraint:

$$\begin{aligned} (3) \quad & \int_0^1 e^{-rt} c(t) dt \\ &= \int_0^R e^{-rt} (y(t) - T(y(t))) dt \\ &+ b(Y(\{y(t)\}_{t=0}^R)), \end{aligned}$$

and use it to obtain the following two conditions describing the individual’s optimal choice of work and retirement, i.e., optimality conditions along the intensive and the extensive labor margins,

$$\begin{aligned} (4) \quad & [1 - T_y(y(t))] \\ &+ e^{rt} b_Y \delta_{y(t)} Y(\{y(t)\}_{t=0}^R)] u'(c(t)) \\ &= v\left(\frac{y(t)}{\varphi(t)}\right) \frac{1}{\varphi(t)} \end{aligned}$$

$$\begin{aligned} (5) \quad & [y(R) - T(y(R))] \\ &+ e^{rR} b_Y \delta_R Y(\{y(t)\}_{t=0}^R)] u'(c(R)) \\ &= v\left(\frac{y(R)}{\varphi(R)}\right) + \eta, \end{aligned}$$

where $\delta_{y(t)} Y$ is the Fréchet derivative of Y with respect to $y(t)$ and $\delta_R Y$ is the Fréchet derivative of Y with respect to R .

To analyze the implications of optimizing the parameters of the PIA benefit function, we proceed in two steps. First, we simulate a benchmark allocation by maximizing the social welfare function (2) subject to the individual budget constraint (3) and the individual optimality conditions (4) and (5) for each individual. Doing this, we keep the benefit function $b(\cdot)$ fixed to a stylized PIA benefit function. Second, we solve the same planning problem allowing the planner to also look for the optimal parameters of the stylized PIA benefit function $b(\cdot)$, as described in the previous section, while restricting the total

cost of benefits, accounting for changes in total tax revenue, to not exceed the benchmark cost.

Computing solutions to these problems presents several challenges stemming from the active extensive margin. A continuous version can rely on exact optimality conditions but needs to accurately approximate several functional derivatives as well as integrals where one of the bounds is a choice variable, R . A discrete version avoids approximation but has to have indicator functions and summations with bounds as choice variables within an efficient nonlinearly-constrained optimization. We proceed with a discrete version modeling the indicator functions via auxiliary binary variables, with corresponding upper bounds constraints, and use branch-and-bound mixed-integer programming.

IV. Simulated Pension Systems

In our simulations, we consider discrete version of the model with the following functional form of $U(\theta)$ for each θ :

$$\sum_{t=1}^N \beta^{t-1} \frac{c(t, \theta)^{1-\sigma} - 1}{1 - \sigma} - \sum_{t=1}^{R(\theta)} \beta^{t-1} \left[\frac{1}{\gamma} \left(\frac{y(t, \theta)}{\varphi(t, \theta)} \right)^\gamma + \eta(\theta) \right],$$

where, as a baseline, we set the risk aversion parameter $\sigma = 2$ and the intensive elasticity parameter $\gamma = 2$, which are both well in their respective ranges estimated in the empirical literature. We use calibrated life cycle productivity profiles $\varphi(t, \theta)$ and fixed costs $\eta(\theta)$ from Shourideh and Troshkin (2012). They use individual earnings and hours data, individual retirement ages, and micro estimates of labor supply elasticity at the extensive margin to calibrate discrete versions of this environment, that take into account the individual effective income taxes estimated with TAXSIM as well as social security taxes and benefits.

We use a simple version with five types, $\theta = 1, \dots, 5$, where each type represents a quintile in the distribution of lifetime earnings, and with $N = 12$. Each period in the model corresponds to five years with $t = 1$ corresponding to ages 20–25, $t = 2$ to 25–30, etc. The individuals experience changes in their productivities over

life until $t = N$, at which point they all expire at the same age of 80.

We assume that $T(\cdot)$ consists of two components: a simplified effective federal income tax at a single rate, taken from the Congressional Budget Office (2005, Table 4A) to be 11.8 percent for 2000, and the old-age part of the OASDI tax, 10.6 percent in 2000.

To simulate our benchmark case, we fix the parameters of $b(\cdot)$ to represent a stylized version of PIA retirement benefits as described above, for the year 2000. It is shown as a dashed line in Figure 1 as a function of annualized AIME. We convert all terms to 2000 dollars and set mean annual income in the model to match that in the data. The mean household income in 2000 in the United States in real 2000 dollars was \$57,045 per year,⁶ indicated in Figure 1 by a solid vertical line. According to the Social Security Administration website, the PIA bend points in 2000 were \$531 and \$3,202 of AIME (\$6,372 and \$38,424 per year). We model AIME as the lifetime average income annualized, times five, to account for the length of a period in the model. The slopes are 0.9, 0.32, and 0.15. The OASDI contribution and benefit base was \$76,200 per year in 2000. We use the bend points from 2000 for all ages.⁷

The contrast to the benchmark is the case where the planner chooses the parameters of the benefit function to maximize social welfare subject to individual optimality and individual budget constraints as described in the previous section.

The optimized benefit function is shown in Figure 1 as a solid line. It has the first bend point at \$1,621 of AIME (\$19,451 per year). The marginal benefit between zero and the first bend point is 1. The second bend point, and consequently the optimized benefit base, are effectively unbounded. That is, the optimal marginal benefit after the first bend point is zero. Within our chosen class of simple instruments, the optimal benefit function is the simplest strictly concave function possible—with two segments.

The optimal pension system based on these simple instruments results in an aggregate welfare gain equivalent to 2.72 percent increase

⁶ Census Bureau publication P60-213: DeNavas-Walt, Cleveland, and Roemer (2001)

⁷ This implies that the cohort in the model turned 62 in 2000.




FIGURE 1. RETIREMENT BENEFIT FUNCTIONS

in per-period consumption for every income quintile at every date. About 70 percent of this gain comes from the change in the marginal rates. In particular, optimizing only the slopes, keeping the base and the bend points fixed at their 2000 levels, gives 1.98 percent increase in consumption. In this intermediate case the slopes are 1, 0.637, and 0 shown as the dotted line in Figure 1.

The welfare gains from the optimization disproportionately favor the lower quintiles (first quintile gaining 5.7 percent, second gaining 1.4 percent in the baseline case) with higher quintiles losing (0.01, 1.7, and 4.7 percent respectively for third, fourth, and fifth quintiles).

To examine the robustness of these gains, we vary the parameters of intensive elasticity and risk aversion and the redistributive motive in the social welfare function. The baseline gains decrease with the Frisch elasticity parameter: changing γ from 2 to 3 reduces the gain from 2.72 to 2.42 percent. The gains increase with risk aversion: from 2.72 percent in the baseline case of $\sigma = 2$ to 3.10 percent with $\sigma = 2.5$ and to 5.09 percent with $\sigma = 3$. The gain also increases with redistributive motives and the highest gain of 10.9 percent is at the extreme of the redistributive motives in the Rawlsian case.

To develop intuition for the sources of these gains, we find it instructive to remove the restriction of a zero intercept. Allowing the intercept to be also chosen optimally significantly increases the baseline gain to 11.65 percent with the intercept of \$38,258 per year and zero marginal benefit rate, shown in Figure 1 by a dash-dotted line.

Intuitively, zero marginal rate induces more productive people to work more years and, as a result, pay additional income tax, which creates additional revenue that allows to give even bigger retirement benefits. The planner looks to redistribute toward lower types and at the same time to avoid giving incentives to higher types to leave the labor force early. The lump sum benefit achieves that, while the optimal zero-intercept benefit function constructs the best possible approximation within the restricted class of functions. Indeed, in the zero-intercept baseline case, the examination of the lifetime benefit amount reveals it to be approximately flat across types whereas in the benchmark it is increasing.

Furthermore, we find that the baseline optimized benefit function induces the bottom quintile to leave the labor force five years earlier and the top four quintiles to stay in the labor force five years longer than in the benchmark, producing a total output gain of 3.06 percent.

To further put the sizable welfare gain of 2.72 percent in perspective, contrast it to the benchmark reform case in Huggett and Parra (2010), that focuses on permanent shocks, as we do here, and finds that optimizing the retirement benefit function produces 0.18 percent gain.⁸ As pointed out above, one key difference between our analyses is that the model here admits an active extensive labor supply margin. In their case, roughly 3/4 of the gain come from redistributing consumption, to more closely reflect utilitarian objective, and the rest comes from the changes in the labor supply, which are allowed only along the intensive margin. We find these sources to produce, respectively, roughly 1/2 and 1/2 of the total gain due to larger gain from the changes in the labor allocation that now also adjusts along the extensive margin.

Finally, a larger gain of 2.72 percent we find here with the extensive margin responses is of the same order of magnitude as the gain of 1.15 percent in Huggett and Parra (2010) when they extend their model to include all types of shocks with private information: permanent, persistent, and temporary.

⁸ See Reform 1 in Table 2 in Huggett and Parra (2010).

REFERENCES

- Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger.** 2009. "Taxing Capital? Not a Bad Idea after All!" *American Economic Review* 99 (1): 25–48.
- Congressional Budget Office (CBO).** 2005. "Historical Effective Federal Tax Rates: 1979–2003." Washington, DC: CBO.
- DeNavas-Walt, Carmen, Robert W. Cleveland, and Marc I. Roemer.** 2001. *Money Income in the United States: 2000*. US Census Bureau, Current Population Reports, P60–213. Washington, DC: Government Printing Office.
- Huggett, Mark, and Juan Carlos Parra.** 2010. "How Well Does the U.S. Social Insurance System Provide Social Insurance?" *Journal of Political Economy* 118 (1): 76–112.
- Shourideh, Ali, and Maxim Troshkin.** 2012. "Providing Efficient Incentives to Work: Retirement Ages and the Pension System." Unpublished.