Carnegie-Rochester Conference Series on Public Policy 50 (1999) 213-259
North-Holland www.elsevier.nl/locate/econbase

The risk-sharing implications of alternative social security arrangements[†]

Kjetil Storesletten[‡]
Stockholm University, Institute for International Economic Studies
Chris I. Telmer
Carnegie Mellon University, GSIA, Pittsburgh, PA 15213[§]
and
Amir Yaron[¶]
Wharton School, University of Pennsylvania

Abstract

An important aspect of the current U.S. social security system is the tradeoff between the risk-sharing it provides and the distortions it imparts on private decisions. We focus on this tradeoff as it applies to labor market risk and capital accumulation. Specifically, we compare the current U.S. system to a particular proposal put forth in 1996 by the federal Advisory Council on Social Security (1996). We also examine the merits of abolishing social security altogether. We find that, absent general equilibrium effects, the risk-sharing benefits of the current system outweigh the distortions associated with either the alternative or a system of privately administered pensions. Once we incorporate equilibrium effects, however, the interaction between the social security system, private-savings decisions, and the means with

* *Correspondence to:* Kjetil Storesletten, Stockholm University, Institute for International Economic Studies.

[†]For helpful comments and suggestions we thank participants at the Carnegie-Rochester Conference, the Penn Macro Lunch Seminar, the Aiyagari Memorial Conference on Dynamic Macroeconomics, and the 1998 NBER Summer Institute. We also thank Andy Abel, Per Krusell, Stan Zin and, in particular, our discussant, Selo İmrohoroğlu.

[‡]kjetil@iies.su.se
[§]telmerc@andrew.cmu.edu
[¶]yarona@savage.wharton.upenn.edu

PII: S0167-2231(99)00028-7

which the government finances its nonpension expenditures results in a significant welfare benefit being associated with either reform or abolition. These welfare gains arise in spite of the fact that we explicitly incorporate the 'social security debt': the social cost of meeting the obligations associated with the current system.

1 Introduction

In December of 1996, the federal Advisory Council on Social Security (1996) outlined three alternatives to the current U.S. social security system. Loosely speaking, each of the proposals can be broken down into reforms geared towards shoring up the anticipated system deficit in the year 2030, and more fundamental reforms, geared towards changing the nature of how the system 'saves' for its participants. In this paper we focus on the latter, with an emphasis on how the proposals affect a fundamental aspect of social security: the tradeoff between the provision of risk-sharing and the distortion of savings incentives.

More specifically, our goal is to run a horse race and understand the fundamentals underlying the outcome. We use a quantitative general equilibrium model to ask which of three pension arrangements are preferable: the current U.S. social security system, the main proposal put forth by the Advisory council, or the abolition of publicly-provided pensions altogether. The winner of this contest will be the alternative which weighs most favorably in terms of the tradeoff between risk-sharing and savings distortions. We focus on four aspects of this tradeoff and quantitatively decompose our results accordingly. On the risk-sharing side, we ask how well the redistributive mechanism inherent in social security provides insurance against uncertain fluctuations in wages. We also examine to what extent the imperfect annuity provided by social security helps insure against 'mortality risk': the risk of outliving one's savings. On the distortions side, we focus on the interaction between social security and distortionary capital income taxation and, more importantly, the effect of social security on aggregate wealth, through the disincentive to save which it imparts on its participants.

The environment in which we ask these questions is a stationary overlapping generations model with productive capital and a large number of heterogeneous agents per generation. Idiosyncratic risk is represented by both an exogenous process for labor productivity as well as probabilistic death. Agents can trade in a single asset—productive capital—and choose to do so for both life-cycle and precautionary motives. Output is produced by an aggregate production technology to which agents rent capital and labor, the latter being supplied inelastically. The absence of a labor-supply decision will, of course, negate what are undoubtedly important distortions associated with social security payroll taxation. While we certainly acknowledge this, we feel that our simple setting has the advantage of providing a sharp focus on the specifics of how taxation related to social security affects capital accumulation.

Into this environment we inject an abstract representation of both the U.S. social security system and the Advisory Council's proposal. Social security is run by a government and financed via a payroll tax and the taxation of social security benefits. This government also undertakes an exogenously-given level of expenditure, financed through income taxes applied to capital and labor. The distortionary effects of payroll taxation therefore travel two channels. First, by taxing (exogenous) labor supply and providing retirement benefits according to a nonlinear benefit function, the system directly affects savings incentives, aggregate capital accumulation, and therefore market-clearing returns. Second, through the effect on capital accumulation, the social security system indirectly affects the government's operating budget (by changing both the capital income-tax base and the labor income-tax base) and therefore has implications for the income-tax rate required to finance nonpension expenditures. This indirect channel and the 'feedback' it generates (e.g., lower capital accumulation implies higher income taxes which further distort capital accumulation) capture, in our opinion, an important aspect of social security finance: the fact that the distortionary effects of fiscal policy and pension policy are tightly intertwined.

A noteworthy aspect of our analysis is the means with which we make welfare comparisons. For the most part, we compare the current U.S. system to some alternative by comparing welfare across alternative steady states, each corresponding to a stationary equilibrium with a different social security system. A critical aspect of these measurements is that they incorporate the so-called 'social security debt': the gross obligations which a pay-as-you-go system has to existing generations (or, equivalently, the debt associated with providing a transfer to the initial old generation). Our feeling is that, absent adjustments for this notion of indebtedness, comparisons across steady states are misleading. For example, consider a comparison of the U.S. status quo with a world featuring 100% privately-provided pensions. The welfare of an agent in the former economy will reflect the fact that, unlike an agent in the latter, they must service the debt associated with providing an unfunded transfer to the initial generations (those who were retired when the system began). Given that we are ultimately interested in risk-sharing and marginal savings incentives, we choose to control for these 'transfer effects' and, via an adjustment to the government's balance sheet, make social indebtedness comparable across steady states.

An alternative interpretation of our 'indebtedness adjustment,' one which adds additional support to its attractiveness, is related to the transition between equilibria featuring different social security systems. For example, consider eliminating the current U.S. system while at the same time honoring all existing obligations. One means of accomplishing this is for the government to issue a bond which raises capital sufficient to finance all existing promises. Were this bond to be refinanced perpetually, the cost of meeting these obligations (which is, again, equivalent to the cost of financing the transfer to the initial generations) would be shared by all future generations. In this context, our approach will involve the government doing just that: perpetually taxing agents to finance a bond with face value equal to the 'social security debt' in the status quo. While it is clear that this approach provides only a partial account of the effects of transition (Auerbach and Kotlikoff (1987), De Nardi, İmrohoroğlu, and Sargent (1998), and others provide an explicit treatment), our feeling is that it has content in this regard.

A problem with examining alternative institutional arrangements by comparing steady states is that many things can change, making it difficult to pin down the essential economics driving the results. This is particularly true in computational economics. In response, we provide a quantitative decomposition of our overall welfare results into components attributable to four economic factors: (i) distortions related to the taxation of capital, (ii) changes in the level of aggregate wealth ('general equilibrium effects'), (iii) the extent to which agents have access to some sort of annuities, and (iv) the degree to which social security allows for the pooling of idiosyncratic labor-market risk. The welfare effect of each of these factors is isolated via a set of computational experiments which, unlike commonly-used comparative statics approaches, are progressive in nature, each stripping away an additional contributing factor.

One final aspect of our analysis which deserves mention is how we measure the extent to which agents have any risk to share in the first place. Our overall question is essentially a quantitative one. We ask, given the various parameters associated with the Advisory Council's proposal, how much consumption a hypothetical agent would pay to live in a world with the reformed pension system. The answer will depend critically on how much risk agents are endowed with, because the fundamental tradeoff they face is one of distortions versus risk-sharing. In an extreme case, one in which we endow our agents with absolutely no idiosyncratic risk, social security will have no positive attributes, but will provide only distortions. With this in mind, we borrow on previous work (Storesletten, Telmer, and Yaron (1998), Storesletten, Telmer, and Yaron (1997)) and attempt to carefully measure labor-market risks using data from the panel Study on Income Dynamics (PSID). By doing so we feel somewhat more confident that our quantitative welfare comparisons are relevant for the U.S. economy.

Our work is related to the existing literature as follows. The social security framework we use builds upon Auerbach and Kotlikoff (1987), Hubbard and Judd (1987), İmrohoroğlu, İmrohoroğlu, and Joines (1995) and, in particular, Huggett and Ventura (1997) who conduct an analysis of the 1986 Boskin plan. Our study is distinct in several respects. Foremost, we focus on the 1996 Advisory Council plans and the effects of a social security system which is partially pay-as-you-go and partially fully-funded. Huggett and Ventura's (1997) focus is primarily on the Boskin proposal of providing an annuity which is closely tied to accumulated contributions, and is financed on a pay-as-you-go basis. Huggett and Ventura (1997) also incorporate a labor-supply decision, thereby making for a richer examination of the overall effects of distortionary payroll taxation. On the other hand, we focus more starkly on the implications for capital accumulation, the incorporation of the 'social security debt,' and try to more precisely decompose the various channels through which social security provides risk-sharing. Our results are, in some cases, quite different. Huggett and Ventura (1997) find that, in moving from the current system to a Boskin-like one, the aggregate effects are quite small, but that poor (rich) households can experience substantial welfare losses (gains). We find similar effects once we distinguish agents by age and income, but find, in some cases, relatively large aggregate welfare gains.

Several other noteworthy papers are Bohn (1997), who looks at aggregate risk-sharing in a Diamond (1965) OLG economy with production, Gertler (1997), who examines the implications for social security of a variety of fiscal policy issues, De Nardi, İmrohoroğlu, and Sargent (1998) and Huang, İmrohoroğlu, and Sargent (1997), who examine transitional dynamics in a much more explicit manner than we, and İmrohoroğlu, İmrohoroğlu, and Joines (1998) who analyze the aggregate effects of having a fixed factor of production in a world with an unfunded social security system.

The remainder of the paper is organized as follows. In the next section we outline our overlapping generations model, endow it with a fairly general social security structure, and formulate its various equilibrium conditions. In Section 3 we provide a brief overview of the current U.S. social security system, discuss the proposals put forth by the federal Advisory Council, and then outline how they fit into our abstraction. Section 4 explicitly discusses what we mean by the term 'social security debt' and how we measure it in our theoretical economies. In Section 5 we demonstrate how we calibrate our models, including a discussion of measurements from the PSID and how they relate to the risk-sharing aspect of social security. Section 6 presents quantitative results and Section 7 concludes.

2 Overlapping generations model

The stationary overlapping generations framework we use builds on the work of Huggett (1996), İmrohoroğlu, İmrohoroğlu, and Joines (1995), and Ríos-Rull (1994), as well as subsequent work in Storesletten, Telmer, and Yaron (1998). The economy is populated by H overlapping generations of agents, where each generation consists of a large number of atomistic, heterogeneous agents. We use $h, h \in \mathcal{H} = \{1, 2, \ldots, H\}$, to index the age cohort to which a particular agent belongs. Agents face both idiosyncratic labor-market risk, described below, as well as mortality risk. There are no aggregate shocks. The unconditional probability of surviving up to age h is denoted by ϕ_h, where $\phi_1 = 1$, and $\xi_h = \phi_h/\phi_{h-1}$, $h = 2, 3 \ldots, H$, denotes the probability of surviving to age h, conditional on being alive at age $h-1$. The fraction of the total population attributable to each age cohort is fixed over time at φ_h and the population grows at rate λ.

Each individual agent is characterized by a preference ordering over consumption distributions, an endowment process, an asset market position, and an entitlement to a particular sequence of social security benefits. Preferences for an unborn agent are represented by,

$$E \sum_{h=1}^{H} \beta^h \phi_h u(c_h), \tag{1}$$

where β denotes the utility discount factor, c_h denotes the consumption of an h-year-old agent, u is the standard twice differentiable, strictly concave utility function, and the expectation is assumed to be conditional on the state of the economy prior to birth.

Agents of ages 1 through Q are workers while agents of ages $Q+1$ through h are retirees. Workers are endowed with an exogenous quantity of labor hours (or, equivalently, productive efficiency units) which they supply inelastically to an aggregate production technology. Labor income is then determined as the product of hours worked and the endogenously-determined wage rate. The exogenous process for hours worked is the means with which we introduce heterogeneity. The i'th agent of some working cohort, h, will provide an amount of labor, $n_{i,h}$ which is governed by the following stochastic process,

$$\log n_{i,h} = \kappa_h + \alpha_i + z_{i,h} + \varepsilon_{i,h}, \tag{2}$$

where the parameters κ_h are used to characterize the cohort-specific cross-sectional distribution of average income, α_i and $\varepsilon_{i,h}$ are *i.i.n.d.* with mean zero and variance σ_α^2 and σ_ε^2, respectively. Finally,

$$z_{i,h} = \rho z_{i,h-1} + \eta_{i,h} \quad , \quad \eta_{i,h} \sim N(0, \sigma_\eta^2). \tag{3}$$

An agent's endowment thus has three random components: a transitory component, $\varepsilon_{i,h}$, a persistent component, $z_{i,h}$, and a 'fixed effect,' α_i, which is realized at birth and stays with an agent for life. Each component will play an important role, both in terms of interpreting microeconomic data as well as affecting the allocations in our theory. For example, the amount of cross-sectional variation in the PSID which we attribute to the fixed effects will dictate the amount of labor-market risk which an agent faces, conditional on being alive. The latter will have important implications for the risk-sharing role played by financial markets and the extent to which social security enhances risk-sharing opportunities. It will also have important implications for the net value which various age-wealth cohorts associate with the social security system. For notational simplicity, we will hereafter omit i, the agent-specific subscripts.

Output in this world is produced by an aggregate technology to which individuals rent their labor services and capital. The competitively-determined wage and capital rental rates are denoted W and R, respectively. The production function takes the form,

$$Y = Zf(K, N), \tag{4}$$

where K and N represent per capita capital and labor, respectively, Y represents per capita output, and Z represents secular growth. Given aggregate consumption, C, and the rate of depreciation on aggregate capital, δ, the law of motion for aggregate capital can be written,

$$K' = Y - C + (1 - \delta)K. \tag{5}$$

Each agent's choice problem is one-dimensional: given knowledge of their idiosyncratic status, they simply choose an amount of assets to accumulate, a_h. Asset holdings are restricted to lie in a set $\mathcal{A}$ which consists of fractions of ownership in the aggregate production technology.

Finally, the model features a government which administers the social security system, consumes a per capita quantity of goods, G, each period, and perpetually refinances a per capita quantity of debt, B, at competitively-determined interest rates. The government finances its consumption and its interest payments via a 100% estate tax (*i.e.*, they fully capture all 'accidental bequests' left behind by those who die prior to age H) as well as proportional taxes on wage and capital income of τ_l and τ_k, respectively. This distinction between tax rates on different sources of income serves, primarily, to uncover the degree to which our results are driven by distortionary capital taxation. The social security system is completely self-financed via the taxation of benefits in addition to payroll taxes levied at rates τ_j, where j is an index corresponding to a number of different retirement accounts which comprise the system. We use the notation τ to represent the vector of tax rates, $[\tau_l \ \tau_k \{\tau_j\}]$. A crucial aspect of all these taxes is that they are determined endogenously, in order to achieve budget balance, and that they constitute a fundamental part of what we will call an equilibrium.

Our motivation for including a government in this manner is twofold. First, as is discussed in detail in the next section, we pay careful attention to the so-called 'social security debt': the implicit cost borne by all future generations of providing the initial old generation with unfunded benefits. The debt, B, which we endow our government with is the means with which we will make meaningful comparisons across economies with varying levels of social security indebtedness. Second, a primary feature of our model is the distortionary effect of social security on capital accumulation (*i.e.*, there is no labor-supply decision to be distorted). Our view is that an important channel in which this takes place is an indirect route, via tax policy. In simple terms, social security affects savings incentives, the resulting changes in the capital stock affect income tax policy, which itself imparts a further distortion on savings.

2.1 *The social security system*

The current U.S. social security system and the various plans for restructuring it all share the idea of taxing agents during their working years while providing benefits during retirement years. The plans differ primarily in terms of the function relating taxation to benefits and the extent to which individuals have control over how their contributions are invested. The plans also differ in the extent to which they represent current period transfers—what we will call a *pay-as-you-go* component—or actual investments on behalf of the contributor—what we will call a *fully-funded* component. With this in mind we formulate a general representation of a social security system which encompasses both components.

The key variable in either system—or some hybrid of both—is accumulated contributions. We use ω_h^p to denote the accumulated contributions of an h-year-old individual towards the pay-as-you-go component of a given system, and ω_h^f to denote the accumulation towards the fully-funded component. The vector ω_h is then defined as $\omega_h = [\omega_h^p \ \omega_h^f]$. These accumulation functions, for $j = p, f$, take the following form.

$$\omega_h^j = \left\{ \begin{array}{ll} \omega_{h-1}^j P_j + \min(\tau_j n_h W, \bar{\omega}^j) & h \le Q \\ \omega_Q^j & h > Q \end{array} \right\}, \tag{6}$$

where τ_j is a payroll tax rate, $\bar{\omega}^j$ denotes the maximum, per-period contribution level, and P_j denotes the return function which applies to accumulated contributions. As an example, the current U.S. system is represented by a scalar accumulation function, ω_h^p, a single payroll tax, τ_p, and a return function, P_p, which equals the average growth rate in aggregate wages. A hybrid system, comprised of both pay-as-you-go and fully-funded components, would append to this an additional accumulation function and tax rate, and a return function, P_f, equal to some market-determined rate of return.

Another important feature shared by most proposals is that benefits are regressive with respect to past contributions. More specifically, benefits are typically bounded from below in some fashion, in order to accommodate a certain minimum benefit level, and are related to accumulated contributions via a concave function. We denote the lower bound as $\underline{b}^j$, and formulate the benefits function, b^j, as,

$$b^j(\omega_h^j) = \left\{ \begin{array}{ll} 0 & h \leq Q \\ \max[\underline{b}^j, d_j(\omega_Q^j)] & h > Q \end{array} \right\}, \tag{7}$$

where d_j is a concave function. For simplicity we will use b_h^j to denote $b^j(\omega_h^j)$, the benefits received by an h-year-old agent from social security arrangement j.

Finally, both the U.S. status quo and the alternative proposal we consider recognize, in one way or another, benefits as taxable income when received after retirement. We use $\tilde{b}_h^j$ to denote the taxable component of social security benefits, and assume that the applicable tax rate is τ_l, that which applies to wage income. Our formulation of $\tilde{b}_h^j$ is stark, especially when considered alongside its complex U.S. counterpart. For our pay-as-you-go systems a fixed fraction of total benefits is taxable at the labor income tax rate. For the fully-funded system, a lump-sum payment received at retirement is not taxable, but all subsequent income accruing to this payment is. In Section 3 we provide a further discussion of benefits taxation and how our treatment captures several salient features of the U.S. status quo and the Advisory Council proposals.

2.2 *Dynamic programming problem*

We can now represent the cross-sectional distribution for our economy as a function μ, defined over an appropriate family of subsets of $S = (\mathcal{H} \times \tilde{\mathcal{Z}} \times \mathcal{A})$, where $\tilde{\mathcal{Z}}$ is the product space containing all possible idiosyncratic shocks (permanent and transitory) and all possible values for ω_h. In words, μ is simply a distribution of agents across ages, idiosyncratic shocks, capital holdings, and social security contribution levels. Because our economy does not feature aggregate shocks, we are able to rely on a cross-sectional law of large numbers to ensure that, in any stationary equilibrium, the function μ is fixed.

Recalling that the market-clearing return on capital and the wage rate are R and W, respectively, the decisions of an agent of age h are governed by the following constraints.

$$\begin{aligned} c_h + a'_{h+1} &\leq a_h R - \tau_k a_h (R-1) + n_h W(1-\tau_l-\tau_p-\tau_f) + \\ & \sum_{j=\{p,f\}} (b_h^j - \tilde{b}_h^j \tau_l) \\ a'_{h+1} &\geq \underline{a} \text{ and } a_{H+1} \geq 0, \end{aligned} \tag{8}$$

where a_h denotes beginning of period asset (or capital) holdings and a'_{h+1} denotes end of period asset holdings. Our timing convention is that savings decisions are made at the end of the current period, and returns are paid the following period at the realized capital rental rate.

Denoting the value function of an agent of age h as V_h, an agent's choice problem can be represented as,

$$V_h(z_h, a_h, \omega_h) = \max_{a'_{h+1}} \left\{ u(c_h) + \beta \frac{\phi_{h+1}}{\phi_h} E[V'_{h+1}(z_{h+1}, a'_{h+1}, \omega'_{h+1})] \right\} \tag{9}$$

subject to equations (8).

2.3 *Equilibrium*

Our definition of an equilibrium will include budget balance on the pay-as-you-go social security account as well as the government expenditure account. This implies that the payroll tax rate, τ_p and one of the income tax rates, τ_k or τ_l, must arise endogenously, and will constitute an important component of the fixed-point problem underlying our solution algorithm. With this in mind, we define an equilibrium to consist of these two endogenously-determined tax rates, market-clearing prices R and W, and a set of cohort specific functions, $\{V_h, a'_{h+1}\}_{h=1}^{H}$, such that,

1. The firm's profit maximization problem is satisfied.

$$R = Zf_1(K, N) - \delta + 1 \tag{10}$$

$$W = Zf_2(K, N) \tag{11}$$

2. Individual optimization problems are satisfied (so that $\{V_h, a'_{h+1}\}_{h=1}^{H}$ satisfy equations (9)).

3. Markets clear and aggregate quantities result from individual decisions,

$$K + B = \int_S (a_h + \omega_h^f) d\mu \tag{12}$$

$$N = \int_S n_h d\mu \tag{13}$$

$$G + K' + \int_S c_h d\mu = ZF(K, N) + (1 - \delta)K \tag{14}$$

4. The government budget constraint is satisfied,

$$G + [R - (1 + \lambda)(1 + g)]B = \int_S (\tau_k a_h [R - 1] + \tau_l n_h W) d\mu + E, \tag{15}$$

where λ is the population growth rate, g represents the secular growth in GNP per capita, and E denotes accidental bequests,

$$E = \int_S \varphi_h \frac{(1 - \xi_{h+1})}{1 + \lambda} a_{h+1}[(R - 1)(1 - \tau_k) + 1] d\mu \tag{16}$$

5. The pay-as-you-go component of the social security system is balanced, period-by-period.

$$\int_S [\tau_p n_h W + \tilde{b}_h^p \tau_l] d\mu = \int_S b^p(\omega_h^p) d\mu \tag{17}$$

Conditions 1 and 2 are standard. Conditions 3, the aggregate resource constraints, make clear that individual holdings of financial wealth (bonds plus capital) plus accumulated, fully-funded contributions, add up to productive capital plus government debt. Condition 4 demonstrates that government budget balance involves choosing the income tax rates, τ_l and τ_k, so that accidental bequests plus income-tax revenues equal the sum of expenditure, G, plus interest payments on the outstanding debt. The term $(1 + \lambda)(1 + g)$ reflects the fact that government debt, B, will be held constant as a fraction of GNP, which grows at rate $(1 + \lambda)(1 + g) - 1$. Finally, condition 5 ensures period-by-period social security budget balance. Note that taxable benefits associated with the pay-as-you-go scheme stay within the system. By construction, contributions corresponding to τ_f are allocated to individual, fully-funded accounts.

3 The specifics of social security

The reform proposals put forth by the federal Advisory Council on Social Security (1996) take the form of three specific alternatives: the *Maintenance of Benefits* (MB) plan, the *Individual Accounts* (IA) plan, and the *Personal Security Accounts* (PSA) plan. The MB plan leaves most of the pay-as-you-go, U.S. status quo in place, focusing instead on changes in taxation and the investment policies governing the social security trust fund, the target being anticipated financing shortfalls in the second third of the next century. The

PSA and IA plans dig deeper, each proposing some hybrid system in which a pay-as-you-go mechanism coexists with some form of a fully-funded system.

Loosely speaking, each of these proposals can be broken down into reforms geared towards shoring-up the anticipated system deficit in the year 2030, and more fundamental reforms, geared towards changing the nature of how the system 'saves' for its participants. The MB plan focuses mainly on the former whereas the IA and PSA plans turn an increasing amount of attention to the latter. Our analysis is primarily concerned with the more fundamental reforms although, as we discuss below, we do incorporate the net cost of moving from one system to another and grandfathering-in the new benefits regime. With this in mind, we restrict further attention to the U.S. status quo, the PSA, and an economy in which social security is abolished altogether. We now turn to a more detailed description of each of these arrangements, followed by a description of how we implement them within the context of our theory.

3.1 *The U.S. status quo (SQ)*

The current U.S. social security plan is a pay-as-you-go system, designed so that retirement benefits are financed directly from payroll taxation of the existing workforce. The current payroll tax is 12.4%, 6.2% of which is paid by the employee. Retirement benefits, in 1993 dollars, are based on a 90% replacement ratio for the first $5,000 of indexed (previous) annual earnings, 32% for the next $25,000, 15% for the next $30,000 and 0 for any amount above $61,750. The marginal tax rate applied to benefits is increasing in both nonpension income as well as the overall benefit level. We omit any further description, however, as the actual tax code is marred with exceptions and special clauses. The retirement age is 65 years and benefits are based on a worker's highest average 35 years of contributions.

Our theoretical implementation of the U.S. status quo, which we denote the 'SQ' economy, is formulated as follows. We work backwards, taking benefits function d_p and the return function P_p as given, and then solving for the payroll tax rate, τ_p, which ensures a balanced budget (the fully-funded tax rate τ_f, as well as anything else with an f subscript, is zero for this economy). The return function P_p, commonly referred to as the 'indexation rate,' is set to the rate of growth in real wages, 1.5%. The benefit function, d_p, is characterized by the four pairs of cutoff points and replacement rates from the U.S. system which, in order to obtain units which are well-defined in our model, we convert to fractions of per capita GNP. For example, the $5,000 level (in 1993 dollars) which determines when the replacement ratio drops from 90% to 32% represents 20% of GNP per capita. Finally, given this system of benefits and the indexation rate, we solve for the payroll tax rate which, given the income tax rates, τ_l and τ_k, equates contributions plus revenue from benefit taxation with total benefits. As we discuss below, the resulting payroll tax is realistic, at $\tau_p = 10.92\%$ for our benchmark economy.

In contrast to the complex means with which taxable benefits are determined in the U.S., we take a simple approach; we treat 25% of pay-as-you-go benefits as being taxable, and tax them at the income tax rate. We find that, in terms of expected tax payments and several other simple measures, this function behaves in a qualitatively similar manner with respect to a more realistic implementation. Note, however, that our approach is less progressive than the actual system, which will tend to understate redistribution in our model.

Finally, we provide for a minimum level of benefits of 24% of GNP per capita, and a maximum level for accumulated contributions (which determine benefits) of 2.47 times GNP per capita. These values correspond to the notation $\underline{b}^p$ and $\bar{\omega}^p$ from Section 2. The former is based on the current U.S. minimum of \$572 per month, expressed in annual terms, in 1993 dollars, and as a fraction of GNP per capita (which is roughly \$25,000). The latter is based on the current annual maximum contribution level of \$61,750, again expressed relative to U.S. GNP per capita. Note that our floor level for benefits is an abstraction relative the actual system, where the floor for an individual who qualifies for social security is essentially determined by the minimum wage rate.

3.2 *Personal security accounts (PSA)*

The PSA plan put forth by the Advisory Council is essentially a hybrid of a pay-as-you-go system and a fully-funded system, with roughly 40% of total contributions going towards the latter. Specifically, the proposal envisions a reallocation of 5 percentage points of the current 6.2% employee contribution tax towards privately-owned retirement accounts. Investment restrictions on these accounts would be minimal and, in contrast to the IA plan, participants would not be required to annuitize their PSA accumulations at retirement. The Council also proposes that any benefits received in lump-sum upon retirement go untaxed, although any income accruing to these funds after retirement becomes taxable at regular rates.

The pay-as-you-go component of the PSA proposal would be financed by the taxation of benefits as well as a 7.4% payroll tax, which is what is left of the status quo 12.4%, after the fully-funded component gets its share. Benefits would be paid out as a lump-sum, irrespective of contributions, of \$410 per month (in 1996 dollars), 100% of which would be treated as taxable income. Finally, the PSA proposal includes a 'Supplemental Tax' of 1.52% until the year 2070 to cover transition costs in moving from the status quo to the new system, something which we ignore.

Our abstract version of the PSA is formulated by setting the fully-funded payroll tax rate, τ_f, to 5% and computing the accumulated, fully-funded contribution accounts, ω_h^f, using the market-determined return function, $P_f = R$. The fully-funded component of benefits is then paid out as a lump-sum to each agent who retires at age 65. The latter assumption is extreme relative to the (vague) guidelines provided by the Advisory Council, but we find it informative in providing a sharp contrast to the partial annuities of the SQ system. It will also have important implications for overall budget balance and our endogenously-determined income-tax rates, because pension-related capital income over the retirement years becomes taxable.

The pay-as-you-go component of our model follows the Advisory Council in providing an annual, lump-sum benefit payment of 17.5% of GNP to each retired agent. This value corresponds to monthly payments of $410, in 1996 dollars. Given these benefits, we solve for a payroll tax rate, τ_p, to ensure budget balance. This value turns out to be 5.13% which, when added to the fully-funded tax rate of 5%, makes for an overall payroll tax of 10.13%.

Finally, taxation of benefits is straightforward and follows the spirit of the Advisory Council's suggestions. The lump-sum, pay-as-you-go benefits are fully taxable at the income tax rate. The fixed payment received from the PSA accounts at age 65 is not subject to tax, but any subsequent capital income is fully taxed at the income rate.

3.3 *Privately-provided pensions (PP)*

An important component of the PSA arrangement is that it represents the privatization of just under half of the publicly-provided pension system. Alternatively, it represents a shift from a system of intratemporal transfers between existing generations, to one in which pension benefits represent actual savings and investment. The limit of this is simply the abolition of the pay-as-you-go social security system altogether, an economy we refer to as the 'PP' economy.

The PP economies we compute are environments with zero payroll taxes, in which agents must provide for retirement consumption themselves. As discussed below, we adjust this economy to have a comparable amount of net indebtedness with respect to the SQ environment. We also subject agents to an income-tax rate sufficient to finance both interest-rate payments on B, the government debt, as well as government expenditure, G.

4 The social security burden

An inescapable aspect of a pay-as-you-go social security system is the implicit liability associated with providing an unfunded transfer for the initial generations who were retired, or part-way through their working years, when the system was introduced. Equivalently, this liability can be thought of as the cost of making good on existing obligations, should the government wish to eliminate the social security system altogether. We denote the per capita value of this obligation as D, and refer to it as the 'social security debt.'

A simple (and well-known) example will make things clear. Consider a deterministic, two-generation world where pay-as-you-go social security contributions are 1 each period, the population grows at rate q, and the market return on capital is $r > q$. Agents work (and contribute) while young and retire when old. By virtue of the fact that there will always be more contributors than retirees, the system can finance benefits of $1+q$ each period, thus providing for a return on contributions of q. In this simple world, $D = 1$; the amount of per capita resources required to liquidate the system w while making good on existing obligations is simply equal to current period contributions. An equivalent interpretation is that each member of the initial old generation received a transfer of 1, and this was financed through government borrowing. The 'tax' of $r - q$ paid by each subsequent old generation is exactly the amount required to keep the level of debt at 1, per capita.

While this example is simplistic, it captures the essence of what we mean by the 'social security debt.' The main difference in our stochastic model with many generations is just more complicated accounting. We proceed in the following manner. Redefine q as the internal rate of return on expected social security contributions. That is, q is the number which results in the following equation being satisfied,

$$\sum_{h=1}^{Q} \phi_h \frac{\tau_p W(1+g)^h E(n_h)}{(1+q)^h} = \sum_{h=Q+1}^{H} \phi_h \frac{E(b_h^p(\omega_Q^p) - \tau_l \tilde{b}_h^p)}{(1+q)^h},$$

where, recall, ϕ_h is the unconditional probability of surviving to age h. The term in the numerator on the left is simply the product of the payroll tax rate and the expected wage bill, where we make explicit the growth rate in wages of g. The numerator on the right is the expected, after-tax benefit received from retirement age Q until death at age H. Recalling that λ is the population growth rate, the solution for q turns out to satisfy $1 + q = (1 + g)(1 + \lambda)$, which is essentially the 'pay-as-you-go' condition suggested by the above example.

Next, we use the internal rate of return, q, to obtain a measure of what the system owes an agent of age h. Define ς_h as the ratio of the average present value, discounted at rate q, of what an h year has contributed to what a retired agent has contributed:

$$\varsigma_h = \frac{\sum_{j=1}^{h} \phi_j \tau_p W E(n_j)(1+g)^j/(1+q)^j}{\sum_{j=1}^{Q} \phi_j \tau_p W E(n_j)(1+g)^j/(1+q)^j}.$$

Given this, we interpret the existing obligation of the system towards the average agent of age h as a promise to pay a retirement annuity equal to ς_h percent of what they would have received under the status quo. The relevance of the 'average agent' in this context is that we are interested in system-wide obligations which, by the cross-sectional law of large numbers, can be expressed in terms of the cross-sectional mean. Note also that the 'survivors premium' is incorporated in these calculations in that the shares, ς_h incorporate survival probabilities.

All that remains is to sum-up the system's obligations towards all age cohorts and discount back to the current period at the market interest rate, r. The latter—the fact that we discount at rate r and not q—is important: it effectively implies that the government is able to finance its obligations at below market rates of return, something intrinsic to the pay-as-you-go system. The sum of obligations, weighted by cohort size, φ_h, is what we label the 'social-security debt', D:

$$\begin{aligned} D &= \sum_{h=1}^{Q} \varphi_h (1+g)^{Q+1-h} \sum_{j=Q+1}^{H} \frac{\phi_j}{\phi_h} \frac{\varsigma_h E(b_j^p(\omega_Q^p) - \tau_l \tilde{b}_j^p)}{(1+r)^{j-h}} \\ &+ \sum_{h=Q+1}^{H} \varphi_h \sum_{j=h}^{H} \frac{\phi_j}{\phi_h} \frac{E(b_j^p(\omega_Q^p) - \tau_l \tilde{b}_j^p)}{(1+r)^{j-h}}. \end{aligned} \tag{18}$$

The first term in this expression is the present discounted value of obligations towards cohorts which are currently working. The second term is analogous, but applies to cohorts which are currently retired. The term $(1+g)^{Q+1-h}$ incorporates the fact that contributions are growing larger over the working years (at the rate of growth in wages, g), but that benefits are held fixed after retirement.

Aside from being of interest in its own right, the social security debt, D, is useful in helping us compare environments with differing levels of social security in a sensible way. Our approach is simply to choose B—the constant level of debt which the government must perpetually refinance—so that each of the economies we study features a comparable level of consolidated indebtedness. The thought experiment we have in mind is as follows. Consider, for example, moving from an economy endowed with the U.S. status quo social security system to one with no system whatsoever. We interpret B as the total amount of borrowing which the government must undertake in order to fulfill the obligations associated with the transition. This level of debt is then rolled-over every period, effectively forcing all future generations to pay the cost of transition (or, equivalently, the cost of the transfer to the initial old). In Section 6, when we incorporate this debt adjustment, we provide more explicit details on how it is computed (*e.g.*, which discount factor, r, is used).

Incorporating the level of indebtedness in this fashion is, in our opinion, an important aspect of our analysis. To not do so would ignore the costs of transition as well as confound the welfare effects of risk-sharing with those of redistribution. In regard to the former—the transitional dynamics—our approach is certainly inferior to one which explicitly models the transition (see, for example, De Nardi, İmrohoroğlu, and Sargent (1998)). It does, however, capture some aspect of moving from one system to another in that the government is (implicitly) prohibited from defaulting on existing obligations.

5 Calibration

The fundamental source of idiosyncratic risk in our economy is the process for labor efficiency units, n_h. It is variation attributable to this process which will be the primary target of the risk-sharing technology afforded by the social security system. A plausible quantitative characterization of n_h is therefore crucial for our question which, essentially, asks how well alternative arrangements provide for risk-sharing otherwise not available via decentralized financial markets. We begin this section by describing how we obtain such a characterization, and then move on to explicitly describe how we calibrate our model.

5.1 *Measuring idiosyncratic risk*

In previous work (Storesletten, Telmer, and Yaron (1998), Storesletten, Telmer, and Yaron (1997)), we have argued that the Panel Study on Income Dynamics (PSID) is an attractive data set for measuring the types of labor-market risks (through the window of the process (2)) faced by a wide cross-section of the U.S. population. If anything, the argument for using the PSID is strengthened in relation to a study of the risk-sharing aspects of social security. It is well-known that the PSID is not representative of the relatively wealthy segment of the U.S. population. However, it also seems reasonable that social security is relatively unimportant for this collection of households. With this in mind we provide a brief summary of our previous findings, supplemented with a parameterization which we feel is better suited for the social security question.

Our departure point is that we choose to use labor-market earnings to calibrate the n_h process, in spite of the fact that it is better interpreted as hours worked. The main reasons for doing so are to avoid measurement issues related to indivisible labor supply, which our model abstracts from, and to allow for the straightforward incorporation of the many types of transfers which will comprise an integral component of our interpretation of a household's 'endowment.' In addition, we verify that the statistical properties of the endogenous process for labor income in our model are very similar to those of hours worked (since our theoretical wage process is relatively stable), thereby providing a sense in which we actually have calibrated theoretical income to PSID income.

More specifically, we define a household's 'endowment' as the combined labor-market earnings of all members, plus any transfers received such as unemployment insurance, workers' compensation, transfers from nonhousehold family members, and so on. We include transfers because our model abstracts from the implicit insurance mechanisms which these payments often represent. That is, we wish to measure the amount of income variation which impinges on household financial decisions *net* of risks which are insured against by programs such as unemployment insurance. Along a similar vein, we study the household as a single unit in order to measure household risk net of things like substitution in labor supply between household members in response to some shock.

Our panel of PSID households is constructed as follows. We use data from the surveys dated 1969 through 1992. Since each survey pertains to household data from the previous year, we refer to the time dimension of our panel as being 1968 through 1991. Beginning with the 1968 survey, we construct a sequence of 22 overlapping panels, each with a three-year time horizon. Each sub-panel (beginning with years 1969 through 1989). consists of data on households who reported strictly positive total household earnings (inclusive of transfers) for that year and the next 2 consecutive years. For example, our 1970 panel is essentially a longitudinal panel on 1663 households over the years 1970, 1971, and 1972. We choose to depart from the more standard longitudinal panel for two primary reasons. First, our sequence of three-year panels contains a sufficient time series dimension so as to allow for the identification of all parameters of interest. Second, the selection of three-year panels mitigates a number of problems which one might associate with a 'flat' longitudinal panel (*e.g.*, a panel with 24 time series observations on each household), such as survivorship bias, a necessarily small cross-sectional sample size, and the fact that average age increases by one year for each survey year.

Two final transformations we apply are to deflate nominal income using the CPI and, in order to incorporate differing family size, to divide total household earnings by the number of household members. The end result is 22 overlapping panels, each with a time dimension of 3 years. The cross-sectional distribution of age is quite stable over each of the panels, with an average (across years) mean and standard deviation of 44.05 and 14.71, and a standard deviation for the means of 1.04. The number of households is substantially larger than would be possible in a longitudinal sample, with a mean and standard deviation of 2019 and 220 observations. Further details on the exact composition of our panel are available in Storesletten, Telmer, and Yaron (1998).

In Table 1, row 1, we reproduce point estimates from our previous paper for the following time series process,

$$y_{it} = g_{it}(y_t) + u_{it} \tag{19}$$

$$u_{it} = z_{it} + \varepsilon_{it} \quad , \quad \varepsilon_{it} \sim N(0, \sigma_\varepsilon^2) \tag{20}$$

$$z_{it} = \rho z_{i,t-1} + \eta_{it} \quad , \quad \eta_{it} \sim N(0, \sigma_\eta^2), \tag{21}$$

where y_{it} is the logarithm of the i'th household's labor-market endowment and $g_{it}(y_t)$ is the portion of y_{it} comprising of aggregate shocks as well as deterministic components of household-specific earnings such as unobservable 'fixed effects' and deterministic variation attributable to household age, education level, and so on. In Storesletten, Telmer, and Yaron (1998) we discuss our particular parameterization of g (which follows closely a number of studies in the labor-market dynamics literature), provide estimates, and discuss how sensitive our results are to alternatives. The first row of Table 1 shows that the autocorrelation coefficient is relatively large, at 0.935, and that the conditional standard deviation of the persistent shock process is roughly 90% larger than that of the transitory shocks.

The process (19) explicitly rules out any agent-specific 'fixed effects,' which one might think of as a household-specific intercept term in the equation describing z_{it}. For the purposes of Storesletten, Telmer, and Yaron (1998) (which asks questions related to asset pricing), such an omission was unimportant, particularly from a theoretical perspective (we did, however, check the sensitivity of our estimates by differencing our data). For risk-sharing social security, however, our feeling is that allowing for an intercept term, which is tantamount to decomposing cross-sectional variation into a stochastic component and a deterministic component, *conditional on being born*, may be crucial. Consider, for example, the welfare benefits of alternative social security arrangements for agents belonging to differing age cohorts. These benefits are likely to be mitigated for all those but the unborn, should a substantial portion of the overall cross-sectional variation in the economy be deterministic, after birth.

We incorporate fixed effects by altering the equation for the shock u_{it} as follows.

$$u_{it} = \alpha_i + z_{i,t} + \epsilon_{it} \quad ; \quad \epsilon_{it} \sim N(0, \sigma_\epsilon^2) \quad , \quad \alpha_i \sim N(0, \sigma_\alpha^2). \tag{22}$$

Table 1:
Idiosyncratic Endowment Process: Parameter Estimates

Description	σ_α^2	ρ	σ_η^2	σ_ε^2
No fixed effects	-	0.935	0.061	0.017
Time Series Moments	0.125	0.935	0.061	0.017
Cross-Sectional Moments	0.326	0.980	0.019	0.005

Entries describe point estimates for the idiosyncratic endowment process described in the text:

$$\begin{aligned} u_{it} &= \alpha_i + z_{it} + \varepsilon_{it} \;\; ; \;\; \varepsilon_{it} \sim N(0, \sigma_\varepsilon^2) \\ z_{it} &= \rho z_{i,t-1} + \eta_{it} \;\; ; \;\; \eta_{it} \sim N(0, \sigma_\eta^2) \;\; , \;\; \alpha_i \sim N(0, \sigma_\alpha^2). \end{aligned}$$

The GMM estimates in the first row are reproduced from Storesletten, Telmer, and Yaron (1998), whereas those in the second and third rows incorporate agent-specific 'fixed effects' in the manner described in the text. Specifically, for the parameters labeled 'Time Series Moments' the values of ρ, σ_η and σ_ε are our GMM estimates and the value of σ_α is chosen so that the average dispersion across age-cohorts matches that of the data (Figure 1). For the parameters labeled 'Cross-Sectional Moments' we choose σ_α^2 to match the initial variance in the data, σ_η^2 to match the slope of the profile (or, equivalently, the end-point), and ρ to match its curvature. The value for σ_ε is chosen to keep the ratio of the variance of the persistent shock to the transitory shock the same as that associated with our GMM estimates.

Next, we 'estimate' the magnitude of σ_α using information on how cross-sectional variation differs by age within our panel. In Figure 1 we report estimates of the age-dependent, cross-sectional variance from the PSID (the solid line), as well as the associated population moments from the process (19), amended with fixed effects as in (22). The dashed-dot line (the lower one) shows how our model without fixed effects matches the pattern in cross-sectional dispersion. We clearly miss a substantial amount of the cross-sectional variation in the data, something which the incorporation of fixed effects will rectify. We take two approaches. First, we simply choose the variance of the distribution from which an agent draws intercept term at birth so that the average, theoretical cross-sectional variance matches that of the data. Values represented by this procedure are reported in Table 1, row 2. The resulting age-profile for cross-sectional variance is represented by the dotted line in Figure 1.

The second method we employ is essentially exactly identified GMM, using age-dependent cross-sectional variances to identify ρ, σ_α, and σ_η (we set the value of σ_ϵ so that the ratio of the variance of the persistent shock to the transitory shock is the same as that associated with our GMM estimates). Loosely speaking, we chose σ_α to match the cross-sectional variation associated with the youngest age-cohort, σ_η, to match the slope required to hit the variation associated with agents just ready to retire (the 60-year olds), and ρ to match the curvature of the age-profile. The values which result are reported in the third row of Table 1 and the theoretical age-profile is represented by the dashed line in Figure 1. Note that the implied value for ρ is substantially higher, at 0.98, which corresponds to the fact that the increase in cross-sectional variance as a cohort age dies out at what appears to be a very slow rate. In addition, the implied fraction of an unborn agent's labor-market risk which is associated with a fixed effect (obtained at birth) increases substantially, something we expect to have an important effect on our welfare calculations.

5.2 *Implementation*

We interpret one period in our model as corresponding to one year of calendar time. The aggregate production technology is Cobb-Douglas:

$$Y = ZK^\theta N^{1-\theta}, \tag{23}$$

Following much of the business-cycle literature, we set θ equal to 0.4 (which corresponds to capital's share of national income being 40%) and allow for a 7.8% annual depreciation rate on the aggregate capital stock. The secular growth rate in GNP per capita, by which we normalize all individual quantities in our model, is chosen to be 1.5% per year.

Figure 1
Cross Sectional Variance by Age

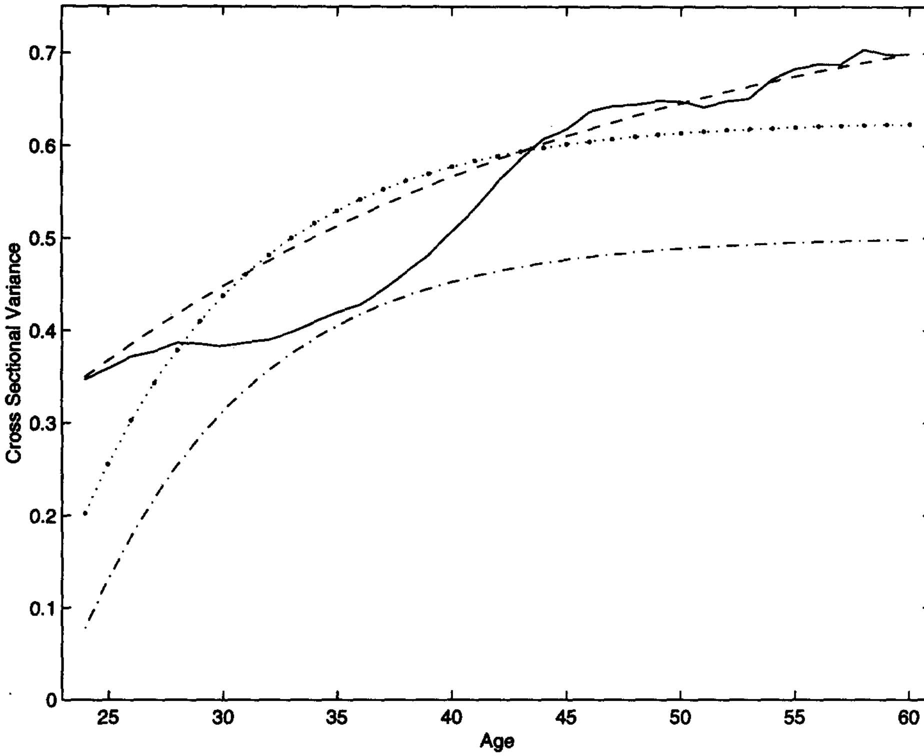

The solid line represents estimates of the cross-sectional variance of PSID labor market income (inclusive of 'transfers'), described in detail in Storesletten, Telmer, and Yaron (1997). The dash-dot line represents population moments associated with the time series process (19), evaluated at parameter estimates obtained by GMM (the first row of Table 1). The dotted line represents the incorporation of 'fixed effects,' where we choose the variance of the distribution from which these parameters are drawn in order to match average dispersion across ages (the second row of Table 1). The dashed line incorporates fixed effects by choosing parameter values in order to match the initial cross-sectional dispersion, and the slope and curvature of the age-profile (the third row of Table 1). Specifically, we choose σ_α^2 to match the initial variance, σ_η^2 to match the slope (or, equivalently, the end-point), and ρ to match the curvature. The value for σ_ε is chosen to keep the ratio of the variance of the persistent shock to the transitory shock the same as that associated with our GMM estimates.

Turning to the characteristics of individual agents, preferences are identical (up to age-dependent mortality risk) and are described by equation (1). We parameterize the period utility function with the standard isoelastic specification,

$$u(c) = \frac{c^{1-\gamma} - 1}{1-\gamma}. \tag{24}$$

We set γ to 2 and the utility discount factor, β, to 1.011 (for details on the interaction between mortality rates and the discount factor, which results in $\beta > 1$, see Hurd (1989)). Demographic variables are chosen to correspond to simple properties of the U.S. work force. Agents are 'born' at age 22, retire at age 65, and are dead by age 100. 'Retirement' is defined as having one's labor income drop to zero. Mortality rates are chosen to match those of the U.S. females in 1991 and population growth is set at 1.0%.

The process for idiosyncratic labor income, equation (2), is implemented as a discrete approximation of the autoregressive time series model and is parameterized using our point estimates from Table 1. In order to highlight the implications of the fixed effects, α_i, we begin by setting them to zero for our benchmark economy. In subsequent experiments (Section 6.4) we allow them to be nonzero and implement them as an *i.i.d.* two-state binomial process, with variance chosen to match our estimates in Table 1. The age-dependent intercept terms, κ_h, are chosen so that the age-dependent mean of the logarithm of labor income in our theory matches our measure from the PSID. The transitory shocks, ε_h, follow a two-state binomial process with equally likely probabilities and a standard deviation of 0.1304. This results in $\varepsilon_h \in \{-0.1304, 0.304\}$. The persistent process is approximated with a 2-state Markov chain.

The only remaining item is the portfolio constraint. In each of the economies we study borrowing is disallowed: $\underline{a} = 0$.

6 Quantitative results

For much of our analysis we will think of an economy endowed with the current U.S. social security system as the benchmark with which to evaluate other alternatives. We treat income from capital and labor equivalently, so that $\tau_l = \tau_k$, and refer to this economy as the 'status quo' (SQ). In Table 2, under the column labelled 'SQ,' we report a number of aggregate features of this economy, to which we now turn.

To begin with, our benchmark economy is broadly consistent with several simple features of the aggregate U.S. economy. The capital to output ratio is 3.10, consumption is approximately half the magnitude of output, and the government consumes just under 20% of aggregate production. Just as importantly, the critical aspects of the social security system line up with the

Table 2:
Properties of Economies with Alternative Social Security Systems
(not adjusted to equalize net indebtedness)

	SQ	PSA	PP
Output, Y	1.636	1.813	2.118
Capital, K	5.067	6.556	9.665
Consumption, C	0.802	0.826	0.810
Government consumption, G	0.311	0.311	0.311
Income tax rate, $\tau_l = \tau_k$ (%)	19.959	18.998	14.090
Payroll tax rate, τ_p (%)	10.923	10.128	
Before tax return, R (%)	5.111	3.262	0.964
After tax return, $R(1-\tau_k)$ (%)	4.091	2.642	0.828
Accidental bequests, E	0.063	0.072	0.119
Pay-as-you-go benefits (net of tax)	0.107	0.056	
Fully-funded benefits		0.064	
Social security debt	2.165	1.515	
Government net worth	-2.165	-1.515	
Publicly-held capital		1.358	
Government-issued bonds, B			
Private financial wealth	5.066	5.199	9.668
IRR on total contributions (%)	2.515	3.209	
IRR on fully-funded contributions (%)		3.262	
IRR on pay-as-you-go contributions (%)	2.515	3.173	
Welfare Losses, relative to SQ (%)			
Proportional (ψ_0)		-7.323	-0.881
Proportional (ψ_1), z_1 = high		-4.562	4.286
Proportional (ψ_1), z_1 = low		-8.933	-3.895
Additive ($\psi_0^{'}$)		-4.233	-0.435
Additive ($\psi_1^{'}$), z_1 = high		-3.771	2.880
Additive ($\psi_1^{'}$), z_1 = low		-4.387	-1.637
Mean Utility (V_0)	-86.460	-80.129	-85.698

Entries correspond to per capita values (except for the various 'rates') for economies featuring the status-quo U.S. social security system (SQ), a system with 100% privately-provided pensions (PP), and the Personal Security Accounts system (PSA). Government debt in these economies *has not* been adjusted to make net indebtedness comparable across steady states. The notation ψ_h denotes the welfare loss, expressed as a proportion of per-period consumption, associated with the average agent of age h moving from the SQ economy into one of the alternatives. The notation $\psi_h^{'}$ is identical, except that the welfare loss is expressed as an additive increment to per-period consumption. IRR denotes 'internal rate of return.' All remaining notations are defined in the text.

U.S. status quo in a reasonable manner. The endogenously-determined rates for income and social security taxes—recall that we fix benefits and solve for tax rates to ensure budget balance—are realistic at roughly 20% and 11%, respectively. The internal rate of return on social security contributions, 2.51%, is substantially smaller than the after-tax return on capital, 4.09%, and matches the growth rate in GNP (recall the discussion in Section 4). Accidental bequests, or what we sometimes refer to as 'estate taxes,' are roughly 3.8% of GNP. Finally, our measure of the gross magnitude of the social security debt (defined in Section 4) is 1.32 times the size of output, a number which is conservative, but in the right ballpark relative to previous studies such as Feldstein (1997) and De Nardi, İmrohoroğlu, and Sargent (1998). These papers explicitly incorporate an anticipated increase in the average age of the population (*i.e.*, the retiring baby-boomers), and, not surprisingly, find that implicit obligations are larger than in our economy, which is characterized by a stationary demographic structure.

It is important to note that we have chosen quantitative values for our model's parameters in order to obtain a realistic *status quo* economy. When we consider alternatives, for instance an economy with only private pensions, it will often be the case that, absent changes in these parameters, aggregate prices and quantities will be unrealistic as a result of fairly dramatic swings in the level of private savings. In these cases, we make a conscious choice not to adjust parameter values, the idea being that holding them fixed results in more direct, meaningful comparisons across economies.

6.1 *Welfare comparisons across alternative economies*

Having described the basic features of our benchmark economy, we now ask how the alternative social security arrangements measure up in terms of welfare as well as a number of aggregate statistics.

Our approach towards obtaining welfare comparisons is standard. We denote the value function for an h-year-old agent living in an economy with the SQ social security system as $V_h(\cdot)$. Similarly, $W_h(\cdot)$ denotes this agent's value function, should they live in an economy endowed with an alternative social security system. The welfare loss associated with moving from the SQ to the alternative is measured as the proportional gain in consumption, received in the alternative, which makes them indifferent between the two. That is, this welfare loss is the number x which results in the following equations being satisfied,

$$V_h(\cdot) = W_h(\cdot; x), \tag{25}$$

where $W_h(\cdot; x)$ solves,

$$W_h(\cdot; x) = max_{a'_{h+1}}\{u(c[1+x]) + \beta\frac{\phi_{h+1}}{\phi_h}EW_{h+1}(\cdot; x)\}, \tag{26}$$

subject to equations (8). The average welfare loss for age-cohort h can now be expressed as the number ψ_h which results in average utility being equated:

$$\int_{\mathcal{Z}\times\mathcal{A}} V_h(\cdot)d\mu = \int_{\mathcal{Z}\times\mathcal{A}} W_h(\cdot;\psi_h)d\mu, \tag{27}$$

where $W_h(\cdot;\psi_h)$ is defined exactly as above, but with the proportional change in consumption held equal across all agents of age h. We also find it instructive to compute these welfare measures in an additive fashion. We denote ψ_h' as the number which results in the previous equations being satisfied, but where the first argument in the maximand is replaced by $u(c+x)$.

What will be crucial about the computation of each of these welfare measures are the prices at which their associated maximization problems are solved. To begin with (in this section of the paper), W_h will be computed vis-à-vis market-clearing prices from the steady state associated with the alternative social security system. That is, V_h and W_h will constitute a comparison *across* different steady states, thereby incorporating the effects of changing equilibrium prices and aggregate quantities into our welfare assessments. In subsequent analysis, when being able to abstract from general equilibrium effects will prove instrumental, W_h will correspond to a problem solved vis-à-vis prices from the reference point equilibrium: the SQ economy. These cases will represent welfare comparisons *within* the benchmark, steady-state equilibrium.

Given a methodology for making welfare assessments, we now turn to our results. To begin with, primarily for illustrative purposes, we compare alternative steady states *without* the adjustment for net indebtedness discussed in Section 4. These comparisons, while useful in a pedagogical sense, are misleading in that they confound risk-sharing with redistribution and ignore any notion of how an economy might transit from one system to another. Accordingly, we conclude this section by incorporating the 'social security debt' and studying steady states in which the initial level of government debt, B, is chosen to make net social indebtedness comparable.

Table 2 reports properties of the stationary steady state corresponding to our three alternative economies: the status quo (SQ), an economy with private pensions (PP), and an economy with the Personal Security Account (PSA) social security system. Our welfare measures indicate that, without an adjustment for the social security system debt, an unborn agent experiences a welfare gain in living in the PP and PSA economies of 0.88 and 7.32%, respectively. Note, however, that not only are these measures misleading in that an indebtedness adjustment has not been made, but also in that the PP economy is dynamically inefficient. In each of our economies population growth is 1% and growth in GNP per capita is 1.5%. GNP therefore grows at 2.5%, which exceeds the PP economy's before-tax return on capital of 0.96%.

The latter is simply a manifestation of the relatively large capital stock—9.66 in the PP world versus 5.07 in the SQ world—which is itself driven by forcing agents to provide for 100% of their own retirement benefits in addition to accumulating a substantially larger precautionary buffer stock. The net result is that, in the PP economy without debt adjustments, an increase in per capita consumption can be attained by simply taxing capital and reducing savings. It is therefore no surprise that the PSA welfare measures indicate a substantially larger gain than their PP counterparts. It is also interesting to note that the main finding of İmrohoroğlu, İmrohoroğlu, and Joines (1995)—that a social security system can help alleviate a dynamic inefficiency problem—is consistent with our result in this context.

We incorporate indebtedness in the following manner. First, as was discussed in Section 4, we obtain a measure of gross obligations in the form of a sequence of promised annuities which correspond to the SQ economy. We then envision the government issuing a bond which raises financing sufficient to meet these obligations. An important question, therefore, involves the appropriate rate of return on capital which will be applied to these government funds. The higher this rate, the lower the requisite amount of borrowing which needs to be undertaken. Were we to explicitly model the transition between steady states, the appropriate rate would be determined within the model. However, since we only compare alternative steady states, the appropriate return is ambiguous. Our approach is to use the return associated with the non-SQ economy: the economy we envision a transition towards. An impressionistic way to think of this is that once the social security system is changed from the status quo, prices and aggregate capital immediately jump to their new steady-state values.

We report results from economies with social security debt adjustments in Table 3. The level of debt associated with the SQ, evaluated at SQ prices, is 2.17, or 1.32 times GNP. In order to finance these obligations, when moving to a system with privately provided pensions, the government needs to borrow 2.95, or 1.60 times GNP. The extra amount here, relative to the debt in the SQ economy, corresponds to the higher capital stock and, therefore, the lower return on capital in the PP economy of 3.02% (versus 5.11% in the SQ). Along similar lines, capital in the PSA economy is slightly larger than in the SQ. The before-tax return on capital is therefore slightly lower, at 4.32%. The requisite amount of borrowing, keeping in mind that only half of the pay-as-you-go obligations are being eliminated, is 1.19, or roughly 70% of GNP.

The welfare implications are quite different once we incorporate the debt adjustment. Table 3 shows that an unborn agent would pay 3.74% of per-period consumption in order to move from the status quo to a debt-adjusted world with only private pensions. They would pay 4.03% in order to move

Table 3:
Properties of Economies with Alternative Social Security Systems
(adjusted to equalize net indebtedness)

	SQ	PSA	PP
Output, Y	1.636	1.706	1.840
Capital, K	5.067	5.631	6.804
Consumption, C	0.802	0.813	0.828
Government consumption, G	0.311	0.311	0.311
Income tax rate, $\tau_l = \tau_k$ (%)	19.959	20.158	13.923
Payroll tax rate, τ_p (%)	10.923	10.054	
Before tax return, R (%)	5.111	4.319	3.018
After tax return, $R(1-\tau_k)$ (%)	4.091	3.448	2.598
Accidental bequests, E	0.063	0.080	0.131
Pay-as-you-go benefits (net of tax)	0.107	0.052	
Fully-funded benefits		0.077	
Social security debt	2.165	1.192	
Government net worth	-2.165	-2.417	-2.954
Publicly-held capital		1.505	
Government-issued bonds, B		1.225	2.954
Private financial wealth	5.066	5.350	9.757
IRR on total contributions (%)	2.515	3.719	
IRR on fully-funded contributions (%)		4.319	
IRR on pay-as-you-go contributions (%)	2.515	3.218	
Welfare Losses, relative to SQ (%)			
Proportional (ψ_0)		-4.027	-3.747
Proportional (ψ_1), z_1 = high		-3.128	-2.545
Proportional (ψ_1), z_1 = low		-4.550	-4.448
Additive (ψ_0')		-2.254	-1.997
Additive (ψ_1'), z_1 = high		-2.584	-1.912
Additive (ψ_1'), z_1 = low		-2.146	-2.027
Mean Utility (V_0)	-86.460	-82.979	-83.221

Entries correspond to per capita values (except for the various 'rates') for economies featuring the status-quo U.S. social security system (SQ), a system with 100% privately-provided pensions (PP), and the Personal Security Accounts system (PSA). Government debt in these economies *has* been adjusted to make net indebtedness comparable across steady states. The notation ψ_h denotes the welfare loss, expressed as a proportion of per-period consumption, associated with the average agent of age h moving from the SQ economy into one of the alternatives. The notation ψ_h' is identical, except that the welfare loss is expressed as an additive increment to per-period consumption. IRR denotes 'internal rate of return.' All remaining notations are defined in the text.

to an economy with a PSA pension system. Note that the former—the value associated with the PP economy—is *larger* than the welfare gain before the debt adjustment. This might seem erroneous when taken at face value, the idea being that welfare should be higher before incorporating the cost of providing an unfunded transfer to the initial old generation. What is going on, however, is that the act of adding societal debt serves to decrease aggregate capital and pull the economy over the dynamically inefficient barrier. Welfare is therefore increased. In the case of the PSA, incorporating the debt adjustment does have the predictable effect of reducing welfare: the gain from switching goes from 7.32% to 4.03% once society is forced to live up to existing obligations.

Exhibit 1: Summary of Experiments Underlying Welfare Decompositions

Notation for Welfare Loss	*Eliminated Effects*	*Comments*
ψ	None	Benchmark welfare comparisons (Table 3).
ψ_a	Capital taxation ($\tau_k = 0$)	Comparison across steady states.
ψ_b	Capital taxation ($\tau_k = 0$) Equilibrium effects	Comparison within steady state.
ψ_c	Capital taxation ($\tau_k = 0$) Equilibrium effects Imperfect annuity effects	Comparison within steady state. Agents given access to perfect annuities.

Our decomposition uses these definitions coupled with the following identity.

$$\underbrace{\psi}_{\text{total welfare loss}} = \underbrace{(\psi - \psi_a)}_{\text{tax effect}} + \underbrace{(\psi_a - \psi_b)}_{\text{equilibrium effect}} + \underbrace{(\psi_b - \psi_c)}_{\text{annuity effect}} + \underbrace{\psi_c}_{\text{risk-sharing effect}} \tag{28}$$

6.2 *Decomposing the welfare gains*

The overall welfare gains associated with social security reform—3.75% of consumption from abolishing the system or 4.03% from privatizing half of it—characterize the outcome of the simple horse race we put forth at the outset of our study. Understanding the economics behind our results is, however, more involved: there are a number of forces at work, each of which can conceivably pull in a different direction. The primary candidates are (i) the interaction between distortionary capital taxation and how the pension system alters the total level of publicly- and privately-held capital, (ii) 'general equilibrium effects': the impact of alternative pension arrangements on the overall level of capital and on market-clearing prices, (iii) the extent to which a social security arrangement provides some form of annuity, thereby aiding in the hedging of mortality risk, and (iv) changes in the income risk-sharing technology available to agents. In order to decompose our overall results into components attributable to each of these effects, we conduct a number of additional experiments, each involving a set of economies in which progressively fewer of the effects are at work. In simple terms, we first eliminate capital taxation effects, then general equilibrium effects, then annuity effects, which finally leaves us with only risk-sharing effects. The differences we find each step of the way constitute the contribution associated with the factor most recently removed. This is, admittedly, a mouthful. The following notation should serve to clarify matters.

The idea behind all this is straightforward. Consider the welfare loss associated with moving from the SQ economy to some alternative, say the PP economy. First, ψ represents the overall welfare loss (from Table 3) which we are attempting to decompose. Each of the four economic factors we have highlighted plays some role in generating this number. Second, ψ_a represents an analogous welfare loss, only computed using economies in which capital income is not taxed. Therefore, of the four original contributing factors, only three remain. The contribution of capital income taxation can thus be associated with the *difference* between the former and the latter: $\psi - \psi_a$. Along a similar vein, if we compute the welfare loss in an environment devoid of both capital-income taxation and equilibrium effects, then the difference between this value and that which excludes only capital taxation constitutes the contribution of equilibrium effects. This difference is labeled $\psi_a - \psi_b$. The story is complete once we eliminate mortality risk, leaving us with the contribution of imperfect annuity effects, $\psi_b - \psi_c$, as well as the residual: the pure contribution of income risk-sharing effects, ψ_c.

The following table provides a quantitative breakdown of how each of the economic factors we focus on contributes towards our bottom line.

The main message of this decomposition is simple. The lion's share of the welfare gain associated with social security reform derives from the general equilibrium effects. As we demonstrate explicitly below, this is a manifestation of the fact that, as a whole, society saves more under either of the PP or PSA arrangements, leading to lower interest rates, a higher capital stock, and a higher level of aggregate output and consumption. Income risk-sharing effects are also important, as evidenced by the 1.6% loss associated with the PP alternative, whereas effects directly attributable to the provision of annuities and capital income-tax distortions play a relatively minor role.

We now provide a more explicit description of the experiments and economic intuition which underlie each of these results.

Exhibit 2: Welfare Decomposition
(welfare losses as a percentage of per-period consumption)

Contributing Factor	PSA	PP
Distortionary capital taxation	0.027	0.717
General equilibrium effects	-3.662	-6.832
Provision of annuities	-0.677	0.758
Income risk-sharing	0.285	1.610
Total	-4.027	-3.747

6.3 *Capital taxation effects*

The economies we use to isolate capital taxation effects are essentially identical to those represented in Table 3, except that all government expenditure is financed through labor-income taxation; $\tau_k = 0$ rather than $\tau_k = \tau_l$ as before. That is, we compute three new steady sates, each corresponding to either the SQ, the PSA, or the PP social security system, where $\tau_k = 0$ in all cases. The welfare numbers of interest, reported along with aggregate statistics in Table 4, are computed by comparing utility in these three steady states, and are thus free of any effects related to capital income taxation.

The main message of Table 4 is that the effects of capital income taxation are not large in magnitude. In the SQ economy, for instance, the increase in steady-state output owing to an elimination of the distortion is a mere 2%. The associated increase in the capital stock is 5%, which generates a decrease in the return on capital from 5.11% to 4.73%. The welfare effects of capital taxation are also not large. Removing the tax implies a small increase in the gain associated with the PSA proposal—from 4.03% up to 4.05%—and a moderate increase associated with private pensions—from 3.75% up to 4.46%. As is outlined above, these changes in welfare gains can be interpreted as the incremental contribution of capital income taxation towards our overall results. We associate a 0.027% effect with the PSA proposal and a 0.717% effect with the PP, both of which are tabulated in Exhibit 2.

6.4 *General equilibrium effects*

What we mean by 'general equilibrium effects' are the implications of changes in our model's institutional structure for the level of aggregate resources and for market-clearing prices. For example, to foreshadow what will turn out

Table 4:
Properties of Economies with Alternative Social Security Systems
(adjusted to equalize net indebtedness, zero capital taxation: $\tau_k = 0$)

	SQ	PSA	PP
Output, Y	1.669	1.729	1.828
Capital, K	5.328	5.822	6.692
Consumption, C	0.802	0.810	0.821
Government consumption, G	0.317	0.317	0.317
Labor income tax rate, τ_l (%)	24.670	24.145	18.440
Payroll tax rate, τ_p (%)	10.788	9.802	
Before tax return, R (%)	4.728	4.078	3.126
After tax return, $R(1-\tau_k)$ (%)	4.728	4.078	3.126
Accidental bequests, E	0.070	0.087	0.133
Pay-as-you-go benefits (net of tax)	0.108	0.050	
Fully-funded benefits		0.074	
Social security debt	2.298	1.190	
Government net worth	-2.298	-2.523	-2.924
Publicly-held capital		1.468	
Government-issued bonds, B		1.333	2.924
Private financial wealth	5.328	5.687	9.616
IRR on total contributions (%)	2.515	3.694	
IRR on fully-funded contributions (%)		4.078	
IRR on pay-as-you-go contributions (%)	2.515	3.377	
Welfare Losses, relative to SQ (%)			
Proportional (ψ_0)		-4.054	-4.464
Proportional (ψ_1), z_1 = high		-3.027	-3.049
Proportional (ψ_1), z_1 = low		-4.645	-5.279
Additive (ψ_0')		-2.198	-2.383
Additive (ψ_1'), z_1 = high		-2.460	-2.333
Additive (ψ_1'), z_1 = low		-2.115	-2.400
Mean Utility (V_0)	-88.307	-84.727	-84.365

Economies are identical to those in Table 3, except that capital income taxation has been abolished: $\tau_k = 0$. Entries correspond to per capita values (except for the various 'rates') for economies featuring the status-quo U.S. social security system (SQ), a system with 100% privately-provided pensions (PP), and the Personal Security Accounts system (PSA). Government debt in these economies *has* been adjusted to make net indebtedness comparable across steady states. The notation ψ_h denotes the welfare loss, expressed as a proportion of per-period consumption, associated with the average agent of age h moving from the SQ economy into one of the alternatives. The notation ψ_h' is identical, except that the welfare loss is expressed as an additive increment to per-period consumption. IRR denotes 'internal rate of return.' All remaining notations are defined in the text.

to be a driving force behind capital accumulation, the existence of annuities turns out to mitigate the desire to save among the working population in our economy. An important implication, one which is external to each individual's decision problem, is that the steady-state level of aggregate capital will be lower than it would be, absent annuities. One might therefore think of the 'externality' associated with annuities markets as the aggregate effect of each agent's decision to save less. That these effects are an important component of what alternative social security systems bring to the table seems obvious. They are certainly implicit in much of the current public-policy debate in the United States, where social security reform is often seen as a remedy for what many feel is an undesirably low savings rate.

In Table 5 we isolate equilibrium effects by holding prices and aggregate quantities fixed across our experiments. The equilibrium effects are then computed as the difference between the fixed-price results and our previous results (Table 4) where endogenous variables were allowed to change.

More specifically, our reference point is an economy endowed with the SQ social security system, in which the capital income-tax rate is set to zero. In this environment, we confront one atomistic agent with an alternative social security system, but hold prices and aggregate quantities identical to those associated with the reference-point equilibrium. In addition, in order to make the comparison meaningful, we tax this atomistic agent so that their net tax burden is comparable to that of an agent facing the SQ system. This is accomplished by computing the present value, in the SQ world, of total payroll, income, and estate taxes paid less total social security benefits received. The atomistic agent, facing one of the alternative systems, is then forced to pay payroll taxes at a rate which implies an equivalent net tax burden.

Table 5 shows the results of these experiments. The first column, labeled SQ, corresponds to an economy identical to that reported in Table 4. The remaining two columns characterize life under an alternative regime, where prices are held identical to those underlying the first column. We see, for example, that the agent forced to live without publicly-provided pensions actually pays a small payroll tax (1.9%) in spite of not receiving any social security benefits. The implication is that the net taxes avoided by opting out of the social security system are not quite offset by the increased estate taxes paid as a result of having to save more for retirement. In regard to welfare, an agent suffers a loss of 2.37% in moving from the SQ to PP economy and a gain of 0.39% in moving to the PSA. Our decomposition methodology therefore associates a welfare gain of 6.83% of consumption with the general equilibrium effects involved in moving from the SQ to the PP economy (tabulated in Exhibit 2). This value is simply the difference between the gain when equilibrium effects are removed, -2.37%, and the gain

Table 5:
Properties of Alternative Social Security Systems, Evaluated at SQ Prices (adjusted to equalize net tax burden, zero capital taxation: $\tau_k = 0$)

	SQ	PSA	PP
Tax burden	-9.810	-9.810	-9.810
Average private financial wealth	5.328	5.708	9.133
Accidental bequests	0.070	0.093	0.138
NPV of cohort consumption	19.046	19.046	19.046
Labor income tax rate, τ_l (%)	24.670	24.670	24.670
Payroll tax rate, τ_p (%)	10.788	11.727	1.900
IRR on total contributions (%)	2.515	3.358	
IRR on fully-funded contributions (%)		4.728	
IRR on pay-as-you-go contributions (%)	2.515	2.342	
Welfare Losses, relative to SQ (%)			
Proportional (ψ_0)		-0.392	2.368
Additive (ψ_0')		-0.204	1.216
Mean Utility V_0	-88.307	-87.961	-90.398

Entries correspond to confronting an agent with alternative social security systems, while holding prices fixed at the level associated with the U.S. status-quo (SQ). PP represents privately-provided pensions, and PSA the personal security account system. The notation ψ_h denotes the welfare loss, expressed as a proportion of per-period consumption, associated with the average agent of age h moving from the SQ economy to some alternative. The notation ψ_h' is identical, except that the welfare loss is expressed as an additive increment to per-period consumption. The tax burden represents the net present value of lifetime taxes paid less benefits received, the measurement of which is discussed in detail in the text. The units in which the tax burden, privately-held capital, and estate taxes are denominated in are essentially per capita, and are directly comparable. IRR denotes 'internal rate of return.' NPV denotes 'net present value.' All remaining notations are defined in the text.

when they are present: 4.46%. Likewise, the welfare gain associated with the PSA is 3.66%, owing to a gain of 4.05% inclusive of equilibrium effects and a gain of 0.39% without them.

The key to understanding these relatively large general equilibrium effects—as well as the main punchline in our paper—is understanding the increase in the capital stock which is associated with making agents provide for their own pensions. In Table 4 we see that this increase is substantial: absent capital income taxation, aggregate capital increases by just under 10% for a move to the PSA system, and just over 25% for a move to the PP. In simple terms, the welfare gains owe to being born into a richer economy. As we will see shortly, and as we alluded to above, a critical aspect of this increase in societal wealth will turn out to be the interaction between savings and the provision of annuities.

6.5 *Imperfect annuity effects*

We isolate effects related to mortality risk by modifying the market for capital to include perfect annuities. This is accomplished through a simple reformulation of the budget constraint, which essentially transfers economy-wide estate taxes from the government back to the private sector. The budget constraint, equation (8), is replaced by,

$$c_h + a'_{h+1}\xi_{h+1} \leq a_h R - \tau_k a_h (R-1) + n_h W(1-\tau_l-\tau_p-\tau_f) + \sum_{j=\{p,f\}} (b^j_h - \tilde{b}^j_h \tau_l). \tag{8'}$$

The only change is that the conditional probability of surviving to age $h+1$, given that one survives to age h, is multiplied against the savings term on the left side of the equation. The idea is that, because one may not survive to capture the benefit of saving, the sacrifice in terms of current consumption is reduced (in an actuarially fair manner). A number of authors, Blanchard and Fischer (1989) for instance, refer to this reduction as the 'survivors premium.'

Table 6 shows what happens when we add perfect annuities to the fixed-price environment with zero capital taxation. Eliminating estate taxation as a means of public finance makes for a substantial increase in the income-tax rate required to balance the budget in the SQ economy. The tax rate increases from 24.7% without annuities (Table 5) to 31.7% with annuities (Table 6). In addition, the adjustment we apply to equate the net tax burden across experiments (recall the above discussion) has the PP agent paying a 5.6% payroll tax—in spite of receiving zero benefits—and the PSA agent paying 12.6%, a slight premium (2%) above the rate applicable in the SQ economy. Again, each of these tax adjustments owes to the fact that the alternative in question involves paying less in net transfers through a reduced size of the pay-as-you-go system.

Finally, the welfare implications of reform, having eliminated capital taxation effects, general equilibrium effects, and annuities effects, are that an agent suffers a 1.6% welfare loss from the abolition of social security and a 0.29% loss from switching to the PSA. Using the methodology which is, hopefully, familiar by now, this translates into a 0.76% loss attributable to changes in annuities markets for the PP economy, and a 0.67% gain associated with the PSA economy.

6.6 *Income risk-sharing effects*

The final economic factor we wish to isolate is labor income risk-sharing. That is, social security is redistributive both in the sense that payroll taxes are proportional and that benefits are regressive with respect to contributions. The floor on benefits can also play a key role, both in our model and the real world.

The good news is that our work here is done. The fixed-price comparisons in Table 6, where capital taxation is eliminated and perfect annuities are added, isolate the risk-sharing effect in an absolute sense, since three of the four candidates we postulated at the outset have been removed. To reiterate, we find that an agent living in a world with private pensions would pay 1.6% of per-period consumption in order to have access to the income risk-sharing technology inherent in the SQ social security system. The analogous number for an agent living in the PSA economy is 0.29%.

An informative context in which to consider the magnitude of these findings is the total amount of idiosyncratic variation faced by agents and how much they would pay to eliminate it. For example, we calculate that in an economy with privately provided pensions, zero income taxation, and no mortality risk (*i.e.*, perfect annuities), an agent would pay roughly 26% of per-period consumption in order to eliminate all idiosyncratic labor-income variation. While this number might seem large, it is actually consistent with previous studies, at least in a methodological sense. Lucas (1987), for instance, calculates that a rough estimate of the welfare gain from eliminating business cycles is given by one-half the product of the risk-aversion coefficient and the variance of aggregate consumption. We use a risk-aversion coefficient of 2, so Lucas's measure suggests that welfare gains from removing variation are roughly on the order of the variance removed. The unconditional variance of consumption in our benchmark economy is approximately 0.16 and 0.32 for agents of age 22 and 65, respectively: numbers which are not unrealistic given evidence from panel data studies (see Deaton and Paxson (1994), for instance). Welfare gains on the order of 26% are therefore plausible, given

Table 6:
Properties of Alternative Social Security Systems, Evaluated at SQ Prices (adjusted to equalize net tax burden, perfect annuities, zero capital taxation: $\tau_k = 0$)

	SQ	PSA	PP
Tax burden	-10.375	-10.375	-10.375
Average private financial wealth	5.218	5.080	7.495
Accidental bequests	–	–	–
NPV of cohort consumption	17.471	17.471	17.471
Labor income tax rate, τ_l (%)	31.667	31.667	31.667
Payroll tax rate, τ_p (%)	10.587	12.596	5.592
IRR on total contributions (%)	2.515	3.175	
IRR on fully-funded contributions (%)		4.889	
IRR on pay-as-you-go contributions (%)	2.515	1.941	
Welfare Losses, relative to SQ (%)			
Proportional (ψ_0)		0.285	1.610
Additive (ψ_0')		0.131	0.749
Mean Utility V_0	-94.128	-94.522	-95.771

Economies are identical to those in Table 5, except that agents have access to perfect annuity markets. Entries correspond to confronting an agent with alternative social security systems, while holding prices fixed at the level associated with the U.S. status quo (SQ). PP represents privately-provided pensions, and PSA the personal security account system. The notation ψ_h denotes the welfare loss, expressed as a proportion of per-period consumption, associated with the average agent of age h moving from the SQ economy to some alternative. The notation ψ_h' is identical, except that the welfare loss is expressed as an additive increment to per-period consumption. The tax burden represents the net present value of lifetime taxes paid less benefits received, the measurement of which is discussed in detail in the text. The units in which the tax burden, privately-held capital, and estate taxes are denominated in are essentially per capita, and are directly comparable. IRR denotes 'internal rate of return.' NPV denotes 'net present value.' All remaining notations are defined in the text.

the amount of idiosyncratic variation we start out with.

Our results on risk-sharing can therefore be thought of in the following loose, but we think informative, manner. An agent would pay 26% of per-period consumption to be able to eliminate idiosyncratic risk. Of these 26 percentage points the risk-sharing component of social security delivers 1.6, or about 6% of the potential gains. The PSA delivers significantly less (just over 1% of the potential gains), reflecting the fact that its fully-funded component represents a movement away from redistribution and towards a stronger link between benefits and accumulated contributions. While these numbers might seem small in an absolute sense, it is important to keep in mind that the overall size of the system is such that only 12% of total wage receipts are incorporated. In light of this, being able to deliver on 6% of the total, by taxing only 12% of the total, seems substantial.

6.7 *The interaction of annuities and aggregate savings*

The upshot of our welfare decomposition is that the lion's share of the gains we attribute to social security reform are due to what we have labeled 'general equilibrium effects': changes in equilibrium prices and aggregate quantities which result from changes in our model's institutional structure. Table 3 demonstrates this in a fairly obvious manner: aggregate capital in the reformed economies increases by 11% in the case of the PSA economy and by 34% in the case of the PP economy. A substantial portion of the associated welfare gain is therefore a simple result of being born into a richer economy. We now demonstrate that a critical ingredient driving this increase in aggregate capital is the savings response of an individual, having lost access to the annuities which social security provides.

In Table 7 we report results from an experiment analogous to that in Table 4, but where perfect annuities have been added in the same manner as described above. Specifically, the comparison in Table 7 is conducted across alternative steady states (with zero capital taxation) in which the need to hedge mortality risk has been eliminated. The idea is that, should we see a smaller increase in aggregate capital relative to our previous experiments, we can attribute much of the aggregate wealth effect in those experiments to the removal of annuities.

Table 7 shows that this is exactly the case. The increase in aggregate capital, vis-à-vis the SQ economy with perfect annuities, is 2.7% for the PSA and 1.9% for the PP. The increases in aggregate consumption and output are also small, especially when compared to the increases reported in Table 4. In addition, a smaller capital stock, relative to the economies which exclude annuities (Table 4), makes for lower wages and therefore a lower payroll tax base. The payroll tax rate must therefore be higher—31.7% in Table 7 versus

Table 7:
Properties of Economies with Alternative Social Security Systems (adjusted to equalize net indebtedness, perfect annuities, zero capital taxation: $\tau_k = 0$)

	SQ	PSA	PP
Output, Y	1.655	1.673	1.667
Capital, K	5.217	5.360	5.315
Consumption, C	0.802	0.804	0.804
Government consumption, G	0.314	0.314	0.314
Labor income tax rate, τ_l (%)	31.667	34.209	36.403
Payroll tax rate, τ_p (%)	10.586	9.165	
Before tax return, R (%)	4.887	4.683	4.746
After tax return, $R(1-\tau_k)$ (%)	4.887	4.683	4.746
Accidental bequests, E			
Pay-as-you-go benefits (net of tax)	0.105	0.042	
Fully-funded benefits		0.083	
Social security debt	2.188	0.914	
Government net worth	-2.188	-2.251	-2.231
Publicly-held capital		1.564	
Government-issued bonds, B		1.337	2.231
Private financial wealth	5.217	5.133	7.546
IRR on total contributions (%)	2.515	4.249	
IRR on fully-funded contributions (%)		4.683	
IRR on pay-as-you-go contributions (%)	2.515	3.817	
Welfare Losses, relative to SQ (%)			
Proportional (ψ_0)		-1.090	0.339
Proportional (ψ_1), z_1 = high		-1.519	-0.997
Proportional (ψ_1), z_1 = low		-0.849	1.086
Additive (ψ_0')		-0.509	0.159
Additive (ψ_1'), z_1 = high		-1.113	-0.716
Additive (ψ_1'), z_1 = low		-0.331	0.429
Mean Utility (V_0)	-94.248	-93.226	-94.567

Economies are identical to those in Table 4, except that agents have access to perfect annuity markets. Entries correspond to per capita values (except for the various 'rates') for economies featuring the status-quo U.S. social security system (SQ), a system with 100% privately-provided pensions (PP), and the Personal Security Accounts system (PSA). Government debt in these economies *has* been adjusted to make net indebtedness comparable across steady states. The notation ψ_h denotes the welfare loss, expressed as a proportion of per-period consumption, associated with the average agent of age h moving from the SQ economy into one of the alternatives. The notation ψ_h' is identical, except that the welfare loss is expressed as an additive increment to per-period consumption. IRR denotes 'internal rate of return.' All remaining notations are defined in the text.

21.7% in Table 4—which itself provides a further drag on disposable income and welfare.

Taken at face value, the implication of these results is that the provision of annuities makes agents worse off. While this is true in a literal sense, we would argue that it confuses the direct effect of annuities—something which our decomposition has shown to be beneficial to agents—with the external effect on aggregate savings. In isolation, an individual agent prefers to have retirement wealth annuitized, thereby avoiding estate taxes and making for a smoother consumption profile over the life cycle. However, the collective implication of annuities is a reduction in aggregate savings, which makes for lower steady-state capital, output, and consumption as well as higher taxes.

Finally, one might think of our results, loosely speaking, as an example of what incomplete markets theorists (see Geanakoplos (1990), for instance) have known for a long time: that adding a market to an incomplete markets setting need not increase welfare. In the incomplete markets literature, many such results are driven by changes in the endogenously-determined space spanned by an agent's budget constraint, the endogeneity arising from relative price effects in multi-good settings. Our example is less rich in two senses. First, our environment features only a single good. Second, in incorporating annuities, we do not really 'add a market,' we simply enhance what the existing market structure is capable of. Nevertheless, the fact that society as a whole can accumulate capital in our setting, and that this directly affects the asset return process, makes 'the span' of our asset markets every bit as endogenous as that in richer models. The resulting externality—the fact that individuals do not incorporate the effect of their savings decision on everyone else's investment opportunities–lies at the heart of the welfare implications we derive.

6.8 *Fixed effects*

Up to this point, the process we have used for idiosyncratic shocks has not included the 'fixed-effect' terms from equation 22, Section 5.1. A simple way of thinking about why these terms might be important involves the way they effectively redistribute total labor-market uncertainty towards the earlier part of the life cycle. To understand this, first consider the process without fixed effects. In this case, uncertainty is spread out uniformly over the earning years (in a conditional sense), a pattern which might be effectively hedged through the *contingent* aspect of buffer stock savings and dissavings (*i.e.*, save upon receiving a good shock, dissave upon receiving a bad one). A redistributive social security system, on the other hand, lacks a strong contingent transfer mechanism and, as we have seen, eliminates (incrementally) only a small fraction of the total idiosyncratic variation faced by agents.

In contrast, consider the case in which fixed effects constitute an important component of total labor-market uncertainty. The ability to make a life-long sequence of state-contingent savings decisions now looks less appealing. Social security, on the other hand, looks more promising as a risk-sharing vehicle; in some sense it represents an uncontingent transfer from rich young agents to old poor agents. One might expect, therefore, to find that the risk-sharing benefits of social security are enhanced in the presence of fixed effects. This will, of course, be especially true for the unborn who are the only cohort for which fixed-effect variation represents uncertainty.

We examine these suppositions in the context of our model by adding fixed effects as is described in Sections 5.1 and 5.2. Specifically, we modify the idiosyncratic risk process to include fixed effects according to the parameter values from the second line of Table 1. We then conduct several experiments designed to isolate pure risk-sharing effects (*i.e.*, Table 6): fixed-price comparisons with perfect annuities and zero capital taxation. The results are reported in Table 8. Under the heading 'Fixed Effects' we report welfare losses from both types of reform, where unborn agents face fixed-effect risk. We also report comparable results from economies with zero-fixed effects, but in which the conditional variance of the persistent shocks has been increased so that the unconditional variance of labor income equals that faced by an unborn agent in the fixed-effect economies. The latter make for meaningful comparisons because the addition of fixed effects serves to add variation in an unconditional sense (see Figure 1).

Our findings confirm the intuition spelled out above. Relative to an economy with no fixed effects and an equal amount of overall variation, an unborn agent suffers a substantially greater welfare loss due to pension reform, should their income process feature a fixed effect. The loss with respect to the PP economy increases from 2.5% to 3.4%. The loss for the PSA experiment increases from 0.6% to 2.1%. Note that, in each case, the welfare loss associated with the economy with no fixed effects is larger than that from Table 6. This is simply a result of increasing the conditional variance in the manner we have described.

Finally, the welfare losses contingent upon age also conform to one's priors. For example, an agent who receives a high fixed-effect shock and a high initial-period persistent shock, gains by almost 1% of consumption from the abolition of public pensions. In contrast, an agent who gets a low fixed-effect shock and a low persistent shock suffers a loss of 6.6% from the same reform.

Our examination of fixed effects, both here and in a related paper (Storesletten, Telmer, and Yaron (1997)), suggests several important points. First, they can provide for a fairly different picture of how much idiosyncratic variation is in the data, and how it is distributed over the life cycle. Second, these differences can have substantial implications for the risk-sharing ben-

Table 8:
Properties of Alternative Social Security Systems, Evaluated at SQ Prices (adjusted to equalize net tax burden, perfect annuities, zero capital taxation: $\tau_k = 0$)

Welfare Losses, relative to SQ	**Fixed Effects**		**High Variance**	
Proportional ψ	PSA	PP	PSA	PP
(ψ_0)	2.107	3.377	0.620	2.511
$(\psi_1, \alpha = \text{low}, \quad z = \text{low})$	3.846	6.580	1.251	3.702
$(\psi_1, \alpha = \text{low}, \quad z = \text{high})$	1.473	2.393	-0.592	0.233
$(\psi_1, \alpha = \text{high}, \quad z = \text{low})$	0.735	0.721	1.251	3.702
$(\psi_1, \alpha = \text{high}, \quad z = \text{high})$	-0.160	-0.937	-0.592	0.233

Welfare losses incorporate a 'fixed effect' into the idiosyncratic income process, as is described in Section 6.4. Entries correspond to confronting an agent with alternative social security systems, while holding prices fixed at the level associated with the U.S. status quo (SQ). PP represents privately-provided pensions, and PSA the personal security account system. Columns labeled 'Fixed Effects' modify the income process according to the second row of Table 1. Columns labeled 'High Variance' set the fixed effect terms to zero, but increase the variance of the persistent shock innovations so as to equate the unconditional variance (from the perspective of an unborn agent) with that of the fixed-effect economies. The notation ψ_h denotes the welfare loss, expressed as a proportion of per-period consumption, associated with the average agent of age h moving from the SQ economy to some alternative. The notation $\psi_1(\alpha, z)$ denotes the welfare loss for an agent of age 1 who received a fixed-effect shock of α and an initial persistent shock of z.

efits associated with both financial markets and publicly-instituted transfer schemes such as social security. Finally, fixed effects are likely to be important for any questions which are age-dependent, for example a more explicit breakdown of the winners and losers in the transition to a new pension system.

7 Conclusions

Our main finding is that agents value the risk-sharing benefits associated with social security, but, once general equilibrium effects are incorporated, not more than the costs associated with its impact on savings incentives. In terms of the federal Advisory Council on Social Security's (1996) proposals, we associate a welfare benefit of 4.02 percent of annual consumption with moving from the status quo to a system of comparable size, but one in which roughly half of worker contributions are earmarked for "personal security accounts (PSA's)': privately-owned, defined-contribution accounts which are invested in capital markets. We associate a slightly lower welfare gain—roughly 3.73 percent of annual consumption—with abolishing social security altogether. A crucial aspect of these welfare comparisons is that they incorporate the 'social security debt': the obligations associated with the status-quo system which we assume the government does not renege upon.

The primary force driving these welfare gains is a kind of externality associated with retirement savings. Social security provides a participant with an imperfect annuity. When that annuity is removed—either completely or partially—individuals save more during their working lives in order to insure against the possibility of outliving their resources during retirement. The collective effect of this increase in savings, something which is external to each individual's choice problem, is an increase in aggregate capital, output, and consumption. This increase in aggregate resources lies at the heart of the welfare gains we uncover.

We argue that this finding—that the provision of annuities can reduce welfare—is not unlike a classic set of results from the literature on general equilibrium with incomplete markets; the endogenous nature of the set of investment opportunities generates an externality which can make changes in the market structure welfare-decreasing. Our example of this is stark in that it abstracts from privately-provided (imperfect) annuities which are, to some extent, available in actual financial markets. Nevertheless, it makes a point which is often overlooked in the debate on social security reform: that the savings response to a change in the system can very much depend on how the availability of annuities is altered. Our results suggest that the quantitative magnitude of this response is substantial.

It is important to note that welfare increases are not necessarily *un fait* *accompli* in our analysis. In spite of its adverse effects on savings incentives, social security plays a valuable role as a risk-sharing technology, both in relation to income risk as well as mortality risk. We quantitatively isolate these effects through experiments which abstract from changing equilibrium prices and aggregate quantities. We find that, for instance, the risk-sharing role of the U.S. status-quo system, when compared to a world without publicly-provided pensions, is worth 2.37 percent of annual consumption from the perspective of an unborn agent. This same agent would pay roughly 34% of annual consumption in order to eliminate idiosyncratic risk altogether, so social security gets them almost 7 percent of the way there. While 7 percent might not seem large, one must keep in mind that it is delivered by taxing just 12 percent of the overall wage bill.

Our analysis does not model the transition between steady states, something which is clearly important in assessing exactly who the winners and losers of any reform would be. In spite of this, one can draw some loose implications of what an explicit account of the transition might yield. First, based on our results which abstract from equilibrium effects, an agent who lives through the early part of the transition, where the level of aggregate capital will be similar to that of the status-quo economy, will suffer slightly as a result of moving to a PSA system and will suffer substantially if publicly-provided pensions are abandoned altogether. These welfare losses are essentially a manifestation of losing access to a valuable risk-sharing technology while not participating in the beneficial aspects of the reformed steady state. It is important to note that this statement is made *net* of the influences of a changing tax burden and/or financing the obligations associated with the status quo. Secondly, based on comparisons which do incorporate equilibrium effects, agents who live through the later part of the transition, or those who live in an economy resembling the new steady state, will benefit substantially from either a privatization of a portion of the portion of the status-quo system, or an abolition of the system altogether. These agents will tend to regret the absence of a societal risk-sharing mechanism, but this regret will be far outweighed by the removal of distortionary capital-income taxation and the higher level of aggregate capital and per capita consumption they will enjoy.

Our overall message is that the benefits or losses associated with social security reform must tradeoff the gains we associate with the new steady state—primarily that agents will be born into a richer economy—with the losses borne by those who live through the transition. This points to the importance of the work of Auerbach and Kotlikoff (1987), De Nardi, İmrohoroğlu, and Sargent (1998), and others who explicitly model the transition and are better able to evaluate this tradeoff. A point which such work might take from our study is that the institutional characteristics of annuities markets are likely to play a quantitatively important role, both in relation to transitional issues as well as to the properties of the reformed steady state.

References

Advisory Council on Social Security, (1996). *Report of the 1994-1996 Advisory Council on Social Security*, Volumes 1 and 2.

Auerbach, A.J. and Kotlikoff, L., (1987). *Dynamic Fiscal Policy*, Cambridge University Press.

Blanchard, O.J. and Fischer, S., (1989). *Lectures on Macroeconomics*, MIT Press, Cambridge, MA.

Bohn, H., (1997). Risk Sharing in a Stochastic Overlapping Generations Economy, UCSB Working Paper.

De Nardi, M., İmrohoroğlu, S., and Sargent, T., (1998). Projected U.S. Demographics and Social Security, Working Paper, University of Chicago.

Deaton, A., and Paxson, C., (1994). Intertemporal Choice and Inequality, *Journal of Political Economy*, **102**: 437-467.

Diamond, P., (1965). National Debt in a Neo-Classical Growth Model, *American Economic Review*, **55**: 1126-1150.

Feldstein, M., (1997). Transition to a Fully Funded Pension System: Five Economic Issues, NBER Working Paper No. 6149.

Geanakoplos, J., (1990). An Introduction to General Equilibrium With Incomplete Asset Markets, *Journal of Mathematical Economics*, **19**: 1-38.

Gertler, M., (1997). Government Debt and Social Security in a Life Cycle Economy, NBER Working Paper No. 6000.

Huang, H., İmrohoroğlu, S., and Sargent, T.J., (1997). Two Computational Experiments to Fund Social Security, *Macroeconomic Dynamics*, **1**: 7-74.

Hubbard, G. and Judd, K., (1987). Social Security and Individual Welfare: Precautionary Savings, Liquidity Constraints and the Payroll Tax, *American Economic Review*, **77**: 630-646.

Huggett, M., (1996). Wealth Distribution in Life-Cycle Economies, *Journal of Monetary Economics*, **38**: 469-494.

Huggett, M. and Ventura, A., (1997). The Distribution Effects of Social Security, Unpublished Manuscript, ITAM.

Hurd, M.A., (1989). Mortality Risks and Bequests, *Econometrica*, **57**: 173-209.

İmrohoroğlu, A., İmrohoroğlu, S., and Joines, D.H., (1995). A Life Cycle Analysis of Social Security, *Economic Theory*, **6**: 83-114.

İmrohoroğlu, A., İmrohoroğlu, S., and Joines, D.H., (1998). Social Security in an Overlapping Generations Economy with Land, Working Paper, Department of Finance, Marshall School of Business, University of South Carolina.

Lucas, R.E., (1987). *Models of Business Cycles,* Basil Blackwell.

Ríos-Rull, J.V., (1994). On the Quantitative Importance of Market Completeness, *Journal of Monetary Economics*, **34**: 463- 496.

Storesletten, K., Telmer, C.I., and Yaron, A., (1997). Consumption and Risk Sharing Over the Life Cycle, GSIA Working Paper No. 1997-E228, Carnegie Mellon University.

Storesletten, K., Telmer, C.I., and Yaron, A., (1998). Asset Pricing with Idiosyncratic Risk and Overlapping Generations, GSIA Working paper No. 1998-E226, Carnegie Mellon University.