OXFORD REVIEW OF ECONOMIC POLICY, VOL. 17, NO. 1

SAVING AND PENSION REFORM IN GENERAL EQUILIBRIUM MODELS

MARIACRISTINA DE NARDI
University of Minnesota
SELAHATTIN İMROHOROĞLU
University of Southern California
THOMAS J. SARGENT
Stanford University[1]

In this paper we study saving and pension reform in the context of general equilibrium models. We discuss some of the reasons for the existence of a social security system. We highlight the key features of the US data that a general equilibrium model of saving and social security should match: the life-cycle profiles of savings, consumption and hours worked, and the aggregate capital–output ratio and saving rate. We describe the structure of a general equilibrium model, and discuss how its parameters are estimated or calibrated from the data. We present the quantitative results from some policy experiments previously computed in the literature, first in the context of stationary demographics and then with an ageing population, and highlight the key policy prescriptions that we learn from these experiments. We conclude by providing some directions for future research.

I. INTRODUCTION: THE METHOD AND THE ISSUES

(i) The Method

This paper describes how economists think about pension reform and savings issues. Economists enforce *coherence* by using a *general equilibrium* model. General equilibrium models incorporate budget constraints and feasibility conditions that 'make things add up'. The models also assume that people are purposeful, which is formalized by assuming that people solve constrained optimization problems. These features enforce consistency among the assumptions about events and possibilities in different markets and at different times. In this paper we review some of the main contributions of the literature, underline what are the relevant issues, and discuss the questions that remain open and for which further research is warranted.

[1] We thank Orazio Attanasio, Marco Bassetto, Marco Cagetti, and Jesus Fernández-Villaverde for their helpful comments.

The original general equilibrium models abstracted from time and uncertainty (they were 'static') and so were not equipped to handle savings and pension issues. They were first created in the nineteenth century by Cournot and Walras to study how cross-effects on prices and quantities would sort themselves out in multiple markets. Important work in the twentieth century has extended general equilibrium models to incorporate time and uncertainty by including trades in commodities to be delivered at different dates and contingent upon different random events. Frank Ramsey (1928), Robert Solow (1956), Paul Samuelson (1958), T. C. Koopmans (1960), David Cass (1965), and Alan Auerbach and Lawrence Kotlikoff (1987) were pioneers in using dynamic general equilibrium models to shed light on various aspects of savings issues.

Dynamic general equilibrium models help to organize ideas about savings and pensions, sometimes with surprising results. For example, demographic and productivity data can take values that make it possible for society as a whole to set up a welfare-improving pyramid scheme in the form of an unfunded social retirement programme. Such schemes are reprehensible when private citizens try to implement them (they are derisively called 'Ponzi schemes'), but as social policy they can nevertheless make society better off under the proper conditions. The case for or against such a social scheme can be understood properly only with a general equilibrium model, because its success depends on how the production and savings decisions of successive generations interact. Many countries have implemented unfunded (or pay-as-you-go) social retirement schemes. Many contemporary proposals for reform call for moving towards more fully funded systems. The empirical factors that determine who might benefit and who might suffer from such reforms can be concisely organized in terms of a general equilibrium model.

(ii) The Policy Issues

We focus mainly on issues in reforming public pension and social insurance systems. Before reforming such institutions, one needs a theory about what they are supposed to accomplish. Such theories have been sharpened by being cast in terms of general equilibrium models.

The basic justification for social savings and insurance programmes is that without them, people would not properly save and insure themselves. Potentially, improper levels of savings and insurance can come either from suboptimal behaviour of individuals or from malfunctioning markets. General equilibrium models are useful guides to thinking about both forms of possible suboptimality and about public policies to improve things.

- If they are myopic or cannot control themselves, households might not save enough for the future or insure themselves enough against bad shocks. Apparent myopia and self-control problems need not indicate peoples' irrationality, but can occur because of the structure of their preferences for goods over time. For myopic individuals and those without self-control, social arrangements that force people to save can make them better off.

- There is an 'opposite' social reason for public intervention in society's savings arrangements, based not on decision-makers' myopia or lack of self-control, but on limitations of markets. It is theoretically possible that society saves *too much* either because many markets for insurance are absent (e.g. annuities) or because of how the demographic structure and total factor productivity move over time. A symptom of this 'capital over-accumulation' outcome is that enough capital has been accumulated to drive the rate of return on capital below the sum of the population growth rate and the rate of improvement of productivity. When this happens, it is possible to 'cure' the capital inefficiency by instituting a pay-as-you-go retirement and insurance system. In such a system, part of retirees' income comes from transfers out of the taxes on working people. Similarly, a tax-and-transfer scheme can partly repair the absence of markets for insuring life-span risk and labour income risks.

This paper follows most of the literature in focusing on the second set of reasons for public interventions in savings and retirement arrangements. But see İmrohoroğlu *et al*. (2000), who study social security in an environment with agents who lack self-control.

(iii) Equivalent Policies

General equilibrium models teach us that some policies that superficially seem different have very similar outcomes. One example is an unfunded social security system in which the current workers are taxed to pay pensions to the current retirees. Government debt can be used to engineer a similar transfer of resources, by rolling the debt over time at the appropriate rate of increase without ever repaying it. To see this, consider the following arrangement: the workers save by accumulating government bonds. The retirees gradually sell their holdings of government bonds to the workers, who also pay taxes to service the government debt. Passing government debt from members of one generation to the next implements cross-generational transfers that resemble those accomplished by an unfunded social retirement system.[2] This means that the funding of a social retirement system cannot be evaluated in isolation, but has to be regarded as part of a comprehensive government fiscal policy package.

(iv) Organization of the Paper

The remainder of this paper is organized as follows. Section II describes some facts about life-time consumption and labour supply profiles that applied dynamic general equilibrium models seek to match. Section III describes the adding-up relationships and the basic structure and free parameters of dynamic general equilibrium models. Section IV summarizes how empirical evidence is used to calibrate the free parameters. Section V briefly describes some important key parameters and specification issues about which little is yet known, but whose resolution impinges on social retirement policy issues. To illustrate the basic methods, section VI describes two numerical policy experiments in some detail. Section VII concludes and points to future research.

II. LIFE-CYCLE BEHAVIOUR OF SAVINGS AND CONSUMPTION

(i) Interpreting the Patterns

In applied general equilibrium models we use inputs that restrict the model in certain ways so that these models generate empirically plausible patterns of aggregate and life-cycle behaviour. Most of these inputs relate to life-cycle profiles that are therefore worth describing. Note that these profiles can be used to deliver estimates of parameters of preferences, technology, and endowments, but they are also required to assess the model in terms of producing observed patterns of behaviour in the data. The other features we might want an applied general equilibrium model to reproduce are aggregate facts, and this motivates their description at the end of this section.

Life-cycle profiles

Important features of the data are the hump-shaped profiles of income, consumption of both durables and non-durable goods, hours worked, and net worth over the life cycle. See Attanasio (1994) for a documentation of the life-cycle profiles (also distinguished by education and other important covariates) of many variables in the USA.

Figure 1 shows an age-consumption profile from İmrohoroğlu *et al*. (1995) using the 1987 Consumer Expenditure Survey (CEX) data. The figure plots mean consumption of all households with a head of the indicated age. The definition of consumption used is total consumption including durables.[3] As we can see from the graph, mean household consumption peaks around age 55.

This hump-shaped feature of the data can be generated from the standard life-cycle model by introducing borrowing constraints (which espe-

[2] However, the two systems can have different implications about risk-sharing across generations. See Huang *et al*. (1997).

[3] Adjustments were made to account for household size and productivity growth. Using the median consumption produces a very similar picture and is about as variable as the mean around late-to-middle ages. This is a reflection of the small sample size for older ages.

Figure 1
Age–Consumption Profile

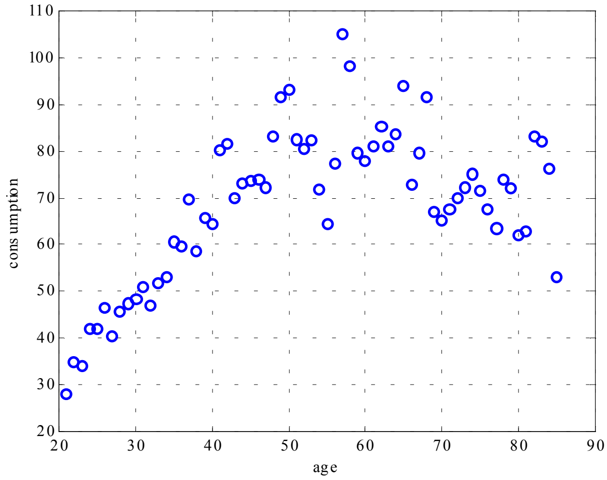

cially prevent the young from consuming as much as they would like) and age-dependent mortality, which effectively makes the individuals discount the future more as they age, thus generating the decreasing consumption profile at later ages. Attanasio *et al*. (1999) have estimated the household's discount factor as a function of family size and female leisure and have shown that these factors can generate much of the hump in the household's consumption profiles, thus providing another reason that we do observe hump-shaped consumption profiles.

Fernández-Villaverde and Krueger (2000) focus not only on reproducing the observed hump-shaped pattern in the consumption of both durables and non-durables, but also want to assess the role of consumer durables in portfolio allocation at different ages (younger households hold a large fraction of their wealth in consumer durables and only later on in life begin to accumulate a sizeable amount of liquid assets). They simulate a model in which durables can be used as a collateral for borrowing and show that this is an important feature to explain household's portfolio composition over the life cycle.

Hurd (1987) studies the path of wealth accumulation over the life cycle to estimate the bequest motive of people with or without children. Fernández-Villaverde and Krueger (2000) also document life-cycle profiles of wealth, reporting the mean and the median net household worth over the life cycle (from the 1995 Survey of Consumer Finances (SCF) household survey data). While both profiles are hump-shaped, there are significant differences between them: (i) the median age–wealth profile is much flatter than the mean age–wealth profile; (ii) both start at positive values at age 25; (iii) there is substantial wealth inequality as the mean profile is always much higher than the median profile; and (iv) both profiles (but especially the mean profile) remain high even after age 80. One feature that appears at odds with a pure life-cycle model is the slow rate at which the elderly decumulate their assets.

In a recent paper, Dynan *et al.* (2000) revisit the life-cycle pattern of saving using data from the Panel Study on Income Dynamics (PSID), the SCF, and the CEX. Median regressions of saving rate on income using consumption as an instrument yield roughly the same outcome across CEX, SCF, and PSID data sets.[4] The saving rate increases with income quintiles, with the effect strongest for the highest income households. Using lagged and future earnings as instruments produces even stronger results for the positive association between the saving rate and predicted income among the working-age households.[5] The age dummy fails to explain the differences in the saving rate. Among older households (ages 62–70), there is still a strong

[4] The measure of saving in the CEX data set is the usual flow definition of income minus consumption, whereas the SCF and the PSID data use the change in wealth as the measure of saving.

[5] The SCF and the PSID data sets produce the strongest correlation.

positive correlation between saving rates and income, with the exception of the SCF which delivers less precise estimates of saving rates. Furthermore, there is little evidence of dissaving among the high-income elderly.

Taking into account differential mortality rates (rich people on average live longer than poorer people), the elderly appear to engage in more dissaving than is revealed by the mean (or median) age–wealth cross-section data (Shorrocks, 1975). Attanasio and Hoynes (2000) use a sample of married couples drawn from the Survey of Income and Program Participation and correct the age–wealth profiles for sample selection due to differential mortality. While this correction is quantitatively important, there remain sizeable amounts of assets that are not annuitized by the elderly. One can think of at least two possible explanations for the pattern of slow decumulation of wealth at advanced ages. One possibility is that households have a bequest motive, and this reduces their decumulation less as they age because of their desire to leave some assets to their descendants. Another explanation can be that uncertainty about out-of-pocket and medical expenditure at advanced ages (Palumbo, 1999) might generate sizeable amounts of precautionary savings for elderly households. Assessing which of these explanations accounts for the bulk of the observed pattern might be relevant also when discussing social security reforms. For example, if people do care about leaving bequests, it would be important to understand exactly how this extra motive for saving affects saving behaviour as this might have very different policy prescriptions. Fuster (1999) studies the effect of eliminating social security on capital accumulation and wealth distribution in a life-cycle model with two-sided altruism. She finds that social security crowds out a much smaller percentage of capital relative to pure life-cycle models. Also, social security in this setting may lead to a more dispersed wealth distribution as the inter-generational redistribution increases the dispersion of bequests.

McGrattan and Rogerson (1998) document age–hours profiles for total population and for various disaggregated cross-sectional time-series and life-cycle profiles, over the last four decades. The age–hours profile in the USA shows a hump-shape, with hours diminishing rapidly after age 55. There are two clusters of discrete drops in hours, first, at age 62, the first age of eligibility for Social Security Benefits, and, second, a larger drop at age 65, the normal retirement age. Retirement age is a key factor in determining the viability of an unfunded social security system when the population is ageing. Incentives are clearly very important in determining when people decide to retire. Gruber and Wise (1999) study social security programmes and early retirement patterns in 11 industrialized countries and show how the provisions of the social security systems themselves typically contribute to withdrawal from the labour force.

(ii) Evolution Over Time

As we mentioned before, the personal saving rate has declined over the last two decades, but the gross saving rate has not changed much. Although there do not appear to be large changes at the aggregate level, there have been significant changes at a more disaggregated level. In a recent paper Krueger and Perri (2000) document the distribution of income and consumption over the last 20 years. Using both the Consumer Population Survey and the CEX data sets, they report Gini coefficients for income and consumption. Their main empirical finding is that income has become significantly more unequal over the last two decades, whereas consumption variability has not changed much. Over the post-war period, the number of weekly hours of market work per person has not changed much. McGrattan and Rogerson (1998) report that this number has stayed at about 22 hours between 1950 and 1990. However, there is a significant change in its composition. The intensive margin, hours per worker, has declined from 40.7 hours per week to about 36.6 hours, whereas the extensive margin, the ratio of employment to population, has increased from 52 per cent to 61 per cent. However, they also report that the constant hours of work coincides with steadily rising real wages over the same period.

McGrattan and Rogerson (1998) also document that the cross-sectional and the life-cycle patterns of labour activity have significantly changed over the last four decades. There have been large decreases of weekly hours of males and larger increases of weekly hours of females. Also, there have been shifts in weekly hours from older to younger and from single-person to married-person households. Their projections, which assume the continuation of current demographic patterns, indicate large changes in age–hours profiles.

(iii) Aggregate Saving and Income

Low rates of personal saving in the USA have been the source of widespread concern among policy-makers. Warnings of 'investment declines' and 'productivity slow-downs' at the aggregate level, and 'inadequate retirement consumption' at the individual level have been given by public speakers. This concern has led policy-makers to search for fiscal policy instruments (401(k)s, Individual Retirement Accounts (IRAs), etc.) to spur private saving. Most of the recent social security reform proposals also aim at raising the private saving rate.

The first panel of Figure 2 shows the US personal saving rate, defined as personal saving as a percentage of disposable personal income. The decline in the personal saving rate over the last 15 years is the source of public concern. Among the reasons cited as driving the decline in personal saving are the increase in household wealth and an expected increase in the growth rate of income. Nevertheless, the decline in the personal saving rate has led some economists to argue that US households do not save 'enough'.[6] Defining saving as net change in assets, which is the definition used by the Federal Reserve Board, one obtains a similar picture.

The private saving rate, which is the ratio of the sum of business saving and personal saving to domestic product, shows much the same time-series behaviour as that of the personal saving rate, but its level is higher, on average, by about 3 percentage points. A more comprehensive measure of saving is national saving, which is the sum of private saving and government saving. The national saving rate shows a decline from 12 per cent in the 1950s and 1960s, to about 6 per cent in the 1980s and early 1990s, although its two components appear to move in opposite directions, especially during the last two decades. Indeed, Velde (1999) argues that the decline in the personal (or private) saving rate seems to coincide with an increase in the government saving. He notes that this evidence is consistent with the permanent-income theory, whereby households view their taxes as future 'investment', allowing 'the government to do their saving for them'.

The second panel of Figure 2 shows the US gross saving rate which appears to have remained roughly constant over the last four decades. Since gross saving is the sum of personal saving, business saving, and government saving, the relative constancy of the gross saving rate reflects increases in business saving (mostly in the form of increases in retained earnings and consumption of fixed capital), and government saving (mostly in the form of smaller budget deficits and depreciation of government fixed capital).

Lusardi, Skinner, and Venti's paper in this issue contains an in-depth discussion of these and other measures of saving, and their relevance for various policy issues.

The next two sections describe how general equilibrium models have been used to extract essential parameters from these patterns for the purpose of studying hypothetical policy experiments, examples of which we describe in later sections.

III. GENERAL EQUILIBRIUM MODELS

(i) An Economy as a Spreadsheet

The general equilibrium models used to study pension reform issues are like big spreadsheets. Let the columns register periods of time, say years, and let the row entries record the decisions or variables impinging on a person or group of similar persons about a particular variable (such as sources of income, taxes, consumption, labour supply, gifts, or bequests to another person). People born earlier appear in earlier rows of the spreadsheet. Each person (or type of person) gets enough rows and

[6] Using the 1983 Survey of Consumer Finances, Avery *et al*. (1984) found that most households held financial assets of less than \$5,000. Even for individuals nearing retirement (ages 55–64), median financial assets were less than \$10,000. Although 75 per cent of households in this age group held positive equity in housing, the median value was \$55,000. There is a large literature that documents saving rates in the United States; for example, see Diamond (1977), Diamond and Hausman (1984), and Hubbard *et al*. (1995), among others.

Figure 2

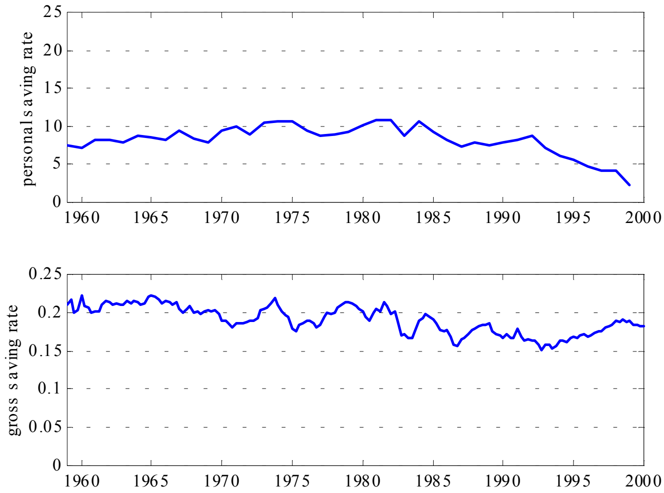

Figure 3
Floor Map

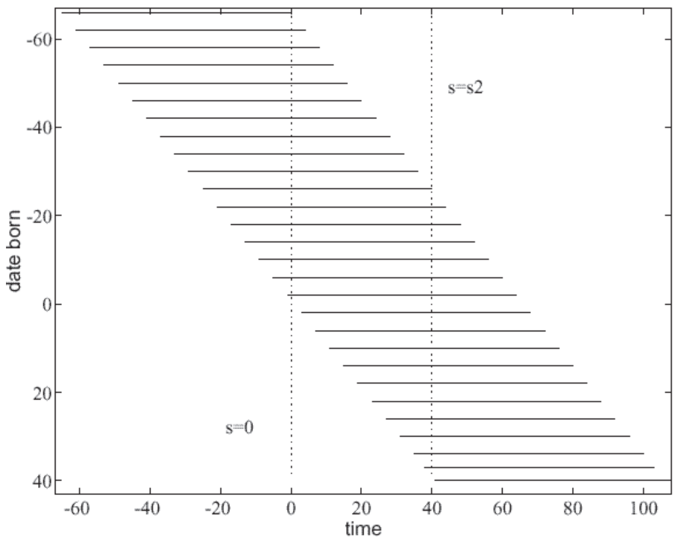

columns to record the important economic decisions they make over their lifetime. Thus, looking across columns at all of an individual's rows tells their lifetime pattern of earnings, tax payments, consumption, and so on. There are as many columns as there are time periods in the economy, thousands of years. A person's column entries differ from zero only during their lifetime.[7] This makes the spreadsheet look striped, with non-zero entries shifting to the right as we move down the spreadsheet because people born later consume and work later. Only governments live forever and so have non-zero entries for all rows. We could put the government at the top of the spreadsheet.

We use Figure 3 to illustrate such a spreadsheet. It is an image of the demography of an overlapping generations model, and will occur as the 'floor' of some three-dimensional graphs that we shall use to record aspects of consumption distributions across agents during some policy transitions. Figure 3 records time along the horizontal axis (corresponding to columns in the spreadsheet) and birth date along the vertical axis (rows of the spreadsheet). A

[7] Because an individual lives for only part of the life of the economy, the rows recording that individual's decisions have non-zero entries only during the lifetime of that individual.

member of a demographic cohort survives from its birthdate to a maximum possible death time of T plus its birthdate, depicted as a *horizontal* line. A vertical line at a given date passes through those cohorts then alive. At each date for each active cohort, for each variable, such as consumption or labour supply, there is a distribution across individuals. These cross-section distributions stand above each pair (birthdate, time) in Figure 3. Later, in a third dimension above the floor, we plot means and standard deviations of some of these distributions (for example, see Figures 4–7). Time s government budget constraints involve sums of distributions at a given point in time, i.e. along vertical lines in Figure 3.

We can use Figure 3 to think about the effects of changes in fiscal and social retirement policy. For example, consider a policy experiment in which everyone born before $-T + s_1$ faces time-invariant tax and benefit rates associated with an initial stationary equilibrium. Assume that everyone born after s_2 faces time-invariant tax and benefit rates associated with a terminal stationary equilibrium. People born between $-T + s_1$ and 0 are 'surprised' at date $s = 0$ by being informed of the transition from s_1 to s_2. They reformulate their life-cycle consumption, saving, and labour supply plans, taking into account the newly announced tax and benefit rates, starting from the asset levels determined by their old saving programmes. Each of the cohorts that live through the transition faces a different sequence of life-time tax and benefit rates. The system settles down to a new stationary equilibrium only after all of the people who have lived through the transition have died off.

The equations comprising the spreadsheet must build in consistency of decisions across rows (i.e. across people) and across columns (i.e. across time). An example of cross-row consistency is that a government transfer to one person must be financed by a tax payment or loan to the government from someone else. Another is that the sum of everyone's labour supply forms the economy's labour input. Across-column consistency is about the inter-temporal consumption and labour supply plans of particular households or types of households. Pensions play an important role in cross-column consistency because households typically consume more than they earn from labour at the end of their life span, and because younger peoples' forecasts of their retirement earnings affect their consumption and labour supply decisions.

The adding-up consistency conditions imposed on such a spreadsheet restrict government policies but do not completely determine them. The period-by-period constraint on government transfers, that they be financed by government receipts from taxes and loans, leaves open many possibilities for using the tax system to transfer resources from one person to another, and from persons of one generation to another. Pension reform is about how to manage such transfers.

(ii) Components

The general equilibrium models used to study social security fill in the spreadsheet by making assumptions about these key components: (i) demographics; (ii) peoples' preferences for goods and leisure over their lifetimes; (iii) production technologies and sources of income; (iv) government fiscal policies; (v) people's expectations about the future; and (vi) an equilibrium concept that reconciles people's decisions with other government policy, other people's decisions, and what is technologically feasible.

- *Demographics*. Because we want to study the consequences of dramatic demographic changes now being forecast by population experts, we take demography as exogenous and deterministic.[8] The demographics are determined by an initial age distribution of population, the rate $n(s)$ of new births at time s, and a sequence of mortality tables $\alpha_t(s)$ that tell the conditional probability at time s that an age t person survives to become an age $t + 1$ person. These objects determine the evolution of the population and its age structure over time, and are perfectly predicted by households. The demographic pattern imparts a structure of gradually overlapping generations.

- *Preferences*. Although households are assumed to be diverse and possibly to live at different

[8] However, for many issues in pension reform in developing countries, it is important to endogenize the birthrate and the mortality probabilities.

times, to conserve on parameters, we often assume that each household shares a utility function (though in equilibrium the arguments vary from household to household because they have different income and wealth levels). The utility function is defined over uncertain streams of consumption and labour supplied c, l. It is usual to assume that an age-t person at time s orders consumption and labour paths according to a time- and state-separable (remaining) life-time utility function of the form

$$E_t(s)\sum_{\tau=t}^{T}\beta^{\tau-t}\left[\prod_{k=t}^{\tau}\alpha_k(s+\tau-t)\right]$$
$$u\left(c_\tau(s+\tau-t), l_\tau(s+\tau-t)\right)$$

where β is a subjective discount factor, T is the maximum survivable age, and $E_t(s)$ is the mathematical expectations operator conditioned on age-t and time-s information, taken with respect to the agent's subjective joint probability distribution over his/her state-contingent life-time consumption and labour-supply plans.

The one-period utility function is typically given by

$$u(c,l) = (c^{\gamma_1}(1-l)^{1-\gamma_1})^{1-\gamma_2}/(1-\gamma_2).$$

The parameter γ_2 governs both the agent's aversion to consumption risk and his/her willingness to substitute consumption and labour inter-temporally over time, whereas the parameter γ_1 governs the worker's static labour–leisure choice.

- *Technology, endowments, and sources of income risk*. An aggregate production technology is typically imported from Solow (1956) to be of the form

$$Y(s) = A(s)K(s)^{\tilde{\alpha}}N(s)^{1-\tilde{\alpha}},$$

where $Y(s)$, $K(s)$, and $N(s)$ are aggregate output, aggregate capital stock, and aggregate labour input, respectively, at time s, $A(s)$ is time-s total factor productivity, and $\tilde{\alpha} \in (0,1)$ is the output share of capital. Aggregate labour supply is taken to be an efficiency and cohort-size weighted average of labour supplied across individuals:

$$N(s) = \sum_i \sum_t l_t^i(s)\varepsilon_t(s)f_t(s)$$

where $\varepsilon_t(s)$ is a time-s relative efficiency index of an age-t worker, i indexes individual workers, and $f_t(s)$ is the time-s population share of age-t individuals.[9] In many models, an exogenous time-invariant age-efficiency profile is posited: $\varepsilon_t(s) = \varepsilon_t$, which is specified to capture the typical age–earnings profile. Furthermore, these models abstract from aggregate uncertainty and assume that the total factor productivity parameter exhibits constant, exogenous growth: $A(s) = (1 + \rho)A(s-1)$.

To provide a framework for analysing income risk, some models add a stochastic component to income or to opportunities to work. One possible specification is that an individual's labour earnings are an exogenous random sequence whose serial correlations resemble those found in earnings from panels of individuals. The volatility and persistence of that income process are important features bearing on retirement policy in models that assume incomplete insurance markets, because those features determine how difficult it is for the individual to use savings to 'self-insure' in the absence of social programmes to insure against adverse income shocks.

- *Fiscal policy*. Fiscal policy is a complete description of a tax system, tax rates, transfers, government portfolio management (including bond sales or purchases), and government purchases of goods and services. Tax rates and transfers affect households directly because they alter the households' budget constraints. The other aspects of fiscal policy can affect households indirectly by affecting the prices and interest rates in their budget constraints.

- *Expectations*. It is usual to impose rational expectations. At each point in time, households make decisions in light of forecasts of future government tax and transfer rates and other factors affecting their budgets. They make each

[9] The cohort shares, $f_t(s)$, reflect time-variation in both the population growth rates $n(s)$, and conditional survival probabilities $\alpha_t(s)$. See De Nardi *et al.* (1999) for details.

forecast using a correct model of the time-series properties of each variable. The only forecast errors that they make are statistically unavoidable.

- *Equilibrium concept and market structure*. The model builder specifies a set of markets, including a description of the times at which and the commodities for which trades can be made. Each traded commodity is assigned a price. An *equilibrium* is a set of prices and quantities such that: (*a*) given the prices, the quantities are such that each household is executing the trades that are in its self-interest, given the household's budget constraint; (*b*) given the prices, firms choose quantities that maximize their profits; and (*c*) supply equals demand on all markets.

IV. CALIBRATION

This section briefly describes how to calibrate general equilibrium models to match some of the data patterns described in section II. Parameter values are chosen to calibrate the model economy so that 'it mimics the world along a carefully specified set of dimensions' (Kydland and Prescott, 1996).

(i) Basic Guidelines for Choosing Parameters

Applied models of social security generally require the specification of four sets of parameters: households, production, government, and demographics. As Browning *et al*. (1999) write:

> For the numerical output of a dynamic equilibrium model to be interesting, the inputs need to be justified as empirically relevant. There are two sources of information that are commonly used in rationalizing parameter values. One is the behaviour of time series averages of levels or ratios of key variables. These time series averages are often matched to the steady-state implications of versions of the models that abstract from uncertainty. The other input is from microeconomic evidence.

Therefore, the basic principle in calibration is to employ econometric estimates of parameters from micro data as much as possible, and for remaining parameters, choose them so that aggregate behaviour in the model economy mimics that of the long-run averages of the macroeconomy. Browning *et al*. (1999) provide a comprehensive discussion of the empirical estimates of these parameters.

There is a large literature on estimating and testing Euler equations generated from alternative economic models to assess the validity of certain model features. Euler equation estimation can also be used to estimate some structural parameters of the model, which are needed for calibration and simulation. Another important finding from this branch of the literature is that the use of disaggregated data appears to provide more support to inter-temporal optimization in the choice of consumption which is an essential feature of economic theory. For example, using the household-level data in the CEX, Attanasio and Weber (1995) argue that the previous rejections of certain economic models are due mostly to the use of aggregate data on consumption, the use of 'food consumption' as a proxy for non-durable consumption, and the inability to account for changes in household composition and the members' labour supply.

We now turn to how these four sets of parameters are chosen in this literature.

Households

For the household sector we have to specify the consumers' preferences and, in models in which the households face idiosyncratic earnings shocks, the stochastic process that governs them.

Different preferences have been used so far in this literature. The household's attitude towards its descendant is a crucial element in determining the impact of an unfunded social security system. Auerbach and Kotlikoff (1987) use a life-cycle framework and find that the capital stock in the US economy could be 24 per cent higher if the social security system were abolished. Barro (1974) shows that if the households behave like dynasties and leave bequests, an unfunded social security has no effect on capital accumulation. While many empirical papers (Hurd, 1987; Hayashi, 1995; Altonji *et al*., 1997) have used micro data sets to test various models of altruistic links across generations, we do not know yet how best to model the links across different generations.[10] Many papers assume that

[10] De Nardi (1999) shows that in the context of a life-cycle calibrated macro model, the introduction of a bequest motive generates a distribution of wealth more consistent with the data.

the households do not care about their offspring. De Nardi *et al*. (1999) study the US unfunded social security system in a set-up in which households care about leaving bequests to their descendants, while Fuster (1999) analyses a model in which there is two-sided altruism between parents and children.

For preferences, many papers assume constant relative risk-aversion utility. Huang *et al*. (1997) study a model with risk-sensitive preferences, a formulation in which the individual's taste for substituting inter-temporally is allowed to differ from their risk aversion (with standard relative risk-aversion preferences, the two parameters are constrained to be the same). The functional form for preferences is dictated by several considerations, such as consistency with balanced growth, eventual consistency with observed aggregate behaviour, and the computational difficulty of solving the model. The last concern arises owing to the presence of certain model features such as individual or aggregate uncertainty, borrowing constraints, and the desire to compute equilibrium transition paths from some initial state to a final steady state. The linear-quadratic preference specification used by Huang *et al*. (1997) and De Nardi *et al*. (1999) allows the use of easier and faster solution methods, thus making it manageable to deal with problems (such as studying demographic transitions) that would otherwise be hard to solve with our current computational resources.

Households have preferences over consumption and leisure at different points in time, and possibly about leaving bequests or about their children's well being. Key parameters entering the utility functions are the discount factor or 'degree of patience', the degree of inter-temporal substitution for consumption and leisure, the degree of intra-temporal substitution between consumption and labour, and the parameters determining 'altruism' towards one's descendants.

The empirical literature that estimates preference parameters based on log-linearized Euler equations does not provide much information on the value of the discount factor, which often is not separately identified. Some recent work has estimated this parameter using structural models (Gourinchas and Parker, 1999; Cagetti, 2000) and micro data sets. Many computational papers in the macroeconomic literature pin down the equilibrium value of the discount factor by using it to match some measure of the capital–output ratio observed in the data.

Several papers have provided micro estimates for the inter-temporal elasticity of consumption (for example, Attanasio and Weber, 1995), and for the inter-temporal elasticity of labour supply (see, for example, Altonji, 1995; French, 2000).

Other parameters, such as the taste for leaving bequests, can be chosen to match macroeconomic aggregates, such as the fraction of total capital transmitted across generations, and microeconomic data to better match the age–wealth profile over the life cycle (De Nardi, 1999).

In models in which the consumers face idiosyncratic earnings shocks, this process also needs to be specified. Based on microeconomic estimates, it is typically assumed that this stochastic component follows an autoregressive process of order one, which means that the realization in one period is a function of the earnings realization in the previous period and an idiosyncratic shock. The two relevant parameters to describe this process are the persistence of the shock and its variance. Many empirical papers based on micro data sets estimate such earnings or income regressions, for instance MaCurdy (1982), Abowd and Card (1989), and Hubbard *et al*. (1995). İmrohoroğlu *et al*. (1999) interpret this idiosyncratic shock as employment opportunities and calibrate it to the aggregate data. Storesletten *et al*. (2000) estimate ranges of parameter values for persistence and variance of the income process from micro data and then use macro models to study which of the parameters values in this range better match macroeconomic observations, such as the fanning out of the variance of income and consumption over the life cycle.

Production

As discussed in section III, in these models the technology is described by an aggregate production function and by the productivity of the household-workers over the life-cycle.

The Cobb–Douglas production function is the most commonly used. The key parameters in this production function are the income shares that go to the various factors of production, such as aggregate (effective) labour, capital, and land. It is standard practice to use aggregate data (such as the National Income and Product Accounts) to pick these parameter values.

The age-profile of the household's productivity is meant to recover the fact that for most people productivity increases at first, as a result of training or human capital accumulation, reaches a peak around mid-career, and then tends to decrease later on. Hansen (1993) and Attanasio *et al*. (1999) are standard references about the deterministic component of household's earnings over the life cycle.[11]

Demographics
In models that allow for life-span uncertainty (starting with Hubbard and Judd (1987) and İmrohoroğlu *et al*. (1995)), the age-dependent survival probabilities are typically taken from survival probability tables. De Nardi *et al*. (1999) study the viability of the current US social security system in an economy in which there is ageing of the population. They use the forecasts of the US Social Security administration for the conditional survival probabilities of the future generations and use a time-varying rate of growth of the population to match the forecast dependency ratio over time (the ratio of people over 65 to workers).

Government
As described in section III, in these models the government levies taxes, finances public expenditure, makes transfers, and services government debt. In the calibration phase, the structure of taxes and transfers is given as a function of the households' actions. One tax rate (usually the tax rate on labour income or the payroll tax) is used to balance the government budget constraint, given the assumptions made on the system of taxes and transfers and on government debt.

The tax (possibly on labour income, capital income, and consumption) and transfer system (social security, and in some cases unemployment) are chosen to match key aspects of the institutional features. Some papers are mostly interested in the inter-generational redistribution implied by the presence of a social security system and therefore abstract from the progressivity of the tax and transfer system by assuming proportional taxes and transfers. Other papers, instead, also study intra-generational redistribution and match some aspects of the progressivity of the taxes and transfer system. Huggett and Ventura (1999) and French (2000) describe and model carefully the structure of social security payments and contributions.

V. IMPORTANT UNKNOWNS

As we choose parameters to match features of the data, some important unknowns remain.

The demographic structure is changing: people live longer and the population is ageing in many countries. This is a key factor in presence of unfunded social security systems, in which the social security benefits to the currently old are paid for by the current workers. In the USA, the Social Security Administration provides demographic projections under three rather different scenarios: 'pessimistic', 'optimistic', and 'normal'. This projection is a key element in assessing the burden of the current system and the costs and benefits of different policy experiments.

The individual's life-span is increasing and people might be working until older ages. We currently do not have much data about the age-efficiency profile at more advanced ages. This is an important variable that influences the individuals' retirement decisions and number of hours worked.

As we previously mentioned, we still do not know how best to model inter-generational altruism. There is a stark difference between what the standard altruistic (dynastic) model of the family predicts about inter-generational transfers of money and the findings of many empirical papers (see, for example, Altonji *et al*., 1997). Villanueva (1999) studies the implications of the altruistic model of the family under imperfect information and moral hazard and

[11] Hansen (1993) uses monthly data derived from the Current Population Survey (published by the Bureau of Labor Statistics in *Employment and Earnings*) and constructs efficiency units as weighted sums of hours worked by each age–sex sub-group in the data, where the weights reflect the relative productivity of that sub-group (average hourly earnings for sub-group divided by average hourly earnings for all sub-groups).

finds that these features account for some of the observed patterns of inter-generational transfers. Inter-generational links are a key ingredient in analysing the effects of social security policies that transfer wealth across generations as they might (partly) undo the effect of social security transfers and taxes across generations.

VI. POLICY EXPERIMENTS

General equilibrium modelling has provided useful insight to many policy questions. In this section, we summarize the outcome of some policy experiments within applied general equilibrium models. We invite the reader to visit the cited papers for the details of the general equilibrium model and its calibration used to generate the quantitative implications described below.[12]

(i) Stationary Demographics

This subsection and the next describe two policy experiments under differing assumptions about demography. This subsection describes a simpler experiment that assumes an unchanging demographic structure. The following subsection then assumes a demographic transition calibrated to forecasts from the US Social Security Administration.

Our first experiment is a fully funded social security system with the sudden elimination of pensions, assuming unchanging demographics. We start from the current unfunded system and surprise all cohorts with the sudden reform. However, we assume that the government fully compensates all people who lose rights to current or future pensions (in the words of G. W. Bush, 'every promise is kept'). In Huang *et al*. (1997), this compensation scheme is tailored so that the value function of the mean person in each 'entitled' cohort is unaltered at the time of the reform. This is accomplished by computing the discounted present value of the social security retirement benefits that the average mean member of each cohort expects to receive at the time of the reform and then making the government give a lump-sum transfer of this amount to each person in the cohort. These payments require a one-time expenditure that the government finances by issuing new one-period debt, called the 'entitlement debt'. The government services and eventually retires that debt by raising the labour income tax rate during some policy transition period.

We can summarize the numerical findings of Huang *et al*. (1997) as follows.

- Eliminating social security raises aggregate capital stock by 40 per cent, and lowers the rate of return on capital from 6.8 per cent to 4.4 per cent.

- The labour income tax rate (which includes the payroll tax for social security in the initial stationary equilibrium) falls from 34 per cent to 14 per cent.

- The 'entitlement debt' is 2.67 times GDP.[13]

Figure 4 shows mean consumption profiles for all of the cohorts affected by the transition. Notice how a version of Figure 3 is laid out on the 'floor' of Figure 4. Aggregate *per-capita* consumption at a point in time s is a weighted-by-time-s cohort size average of the mean consumption of persons of age t and birthdate $s - t$, which corresponds to summing elements corresponding to those associated with a fixed date (a line parallel to the date-born axis in Figure 4, and a line on the diagonal of the floor of Figure 6). Since the sudden termination of social security is accompanied by a complete buy-out of existing cohorts, the age-consumption profiles do not show any changes for existing cohorts, but eventually shift up as the benefits from a larger capital stock trickle down in the form of higher aggregate consumption to all future cohorts.

[12] In particular, both sets of experiments described in this section rely on the computational speed advantage offered by the use of quadratic preferences. Existing formulas from linear-quadratic control theory make it feasible/convenient to compute very detailed policy functions that depend on age and time, and to keep track of a large number of objects along an equilibrium transition path.

[13] In Huang *et al*. (1997), an alternative funding experiment had the government acquire claims on physical capital so that social security benefits can be financed by the returns from publicly held private capital. This scheme for fully funding social security created similar effects on existing cohorts, apparently owing to the quantitative near-equivalence of the debt versus tax-transfer policies in overlapping generations models. However, the benefits to later generations seem to be larger, owing to the maintenance of the annuity role of public pensions.

Figure 4
Evolution of Age–Consumption Profiles

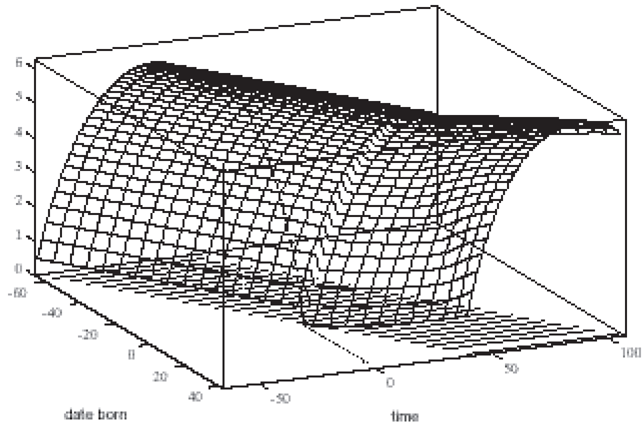

Figure 5
Evolution of Age–Consumption Inequality Profiles

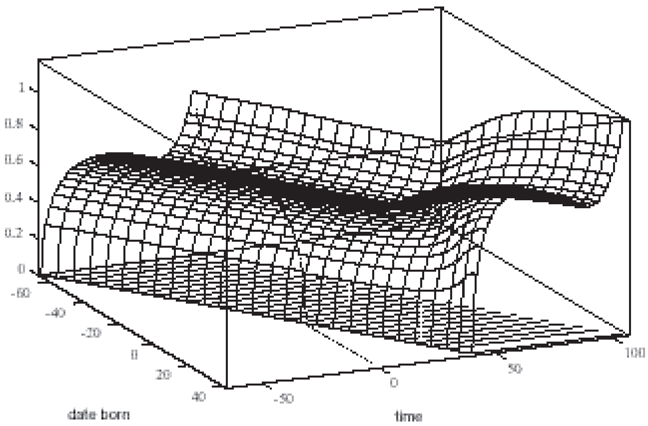

Figure 6
Evolution of Age–Consumption Profiles

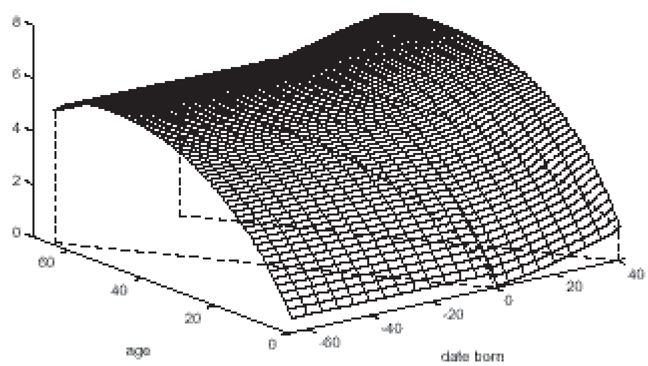

Figure 7
Evolution of Age–Consumption Inequality Profiles

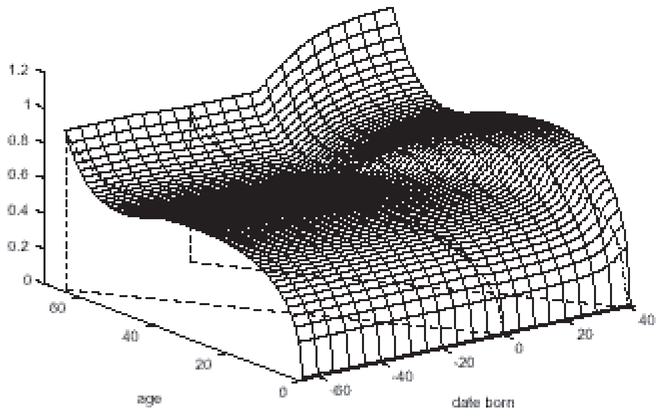

Figure 8
Projected Dependency Ratio

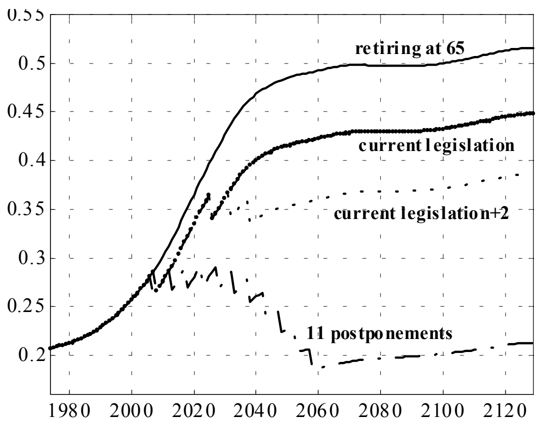

Figure 5 shows how the (cross-sectional) standard deviations of consumption are affected by the transition. Note that social security insures against both mortality risk and idiosyncratic income risk. As this insurance device is removed and not replaced by private markets that are assumed to be missing, consumption inequality worsens among the members of a given cohort as that cohort ages (for a given date, the life-cycle profile of the standard deviation of consumption rises faster than earlier dates).

Figures 6 and 7 represent the same information in Figures 4 and 5, respectively, in a more compact but harder to read way, by plotting a mean or standard deviation against *date born* on the *x*-axis and *age* on the *y*-axis. With these axes, a given calendar date is a diagonal along the 'floor' of the graph, and economy-wide averages must be calculated by computing weighted sums along this diagonal.

(ii) The Ageing of the Population

The US population projections by the Social Security Administration portray a gloomy picture of social security's future: compared with one retiree's pension being financed by five workers, a dependency ratio of 20 per cent in 1990, there will be only two-and-a-half workers to pay for the pension of a retiree in 2060, according to current legislation.[14] In

[14] The demographic projections for Japan, and some of the European countries such as Italy and Germany, are even gloomier in this respect.

Figure 9
Welfare Effects Relative to Experiment 1
(compensation for each generation normalized by their mean assets at birth)

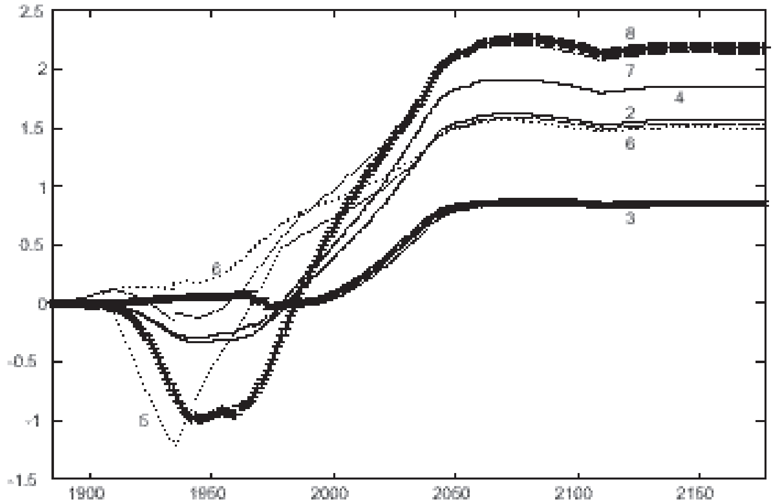

Figure 8, De Nardi *et al*. (1999) document four projected dependency ratios that correspond to alternative eligibility rules: (i) continuing the current age-65 qualification; (ii) adhering to the current legislation that postpones the normal retirement age to 66 in 2008 and to age 67 in 2026; (iii) adding two additional postponements beyond these two; and (iv) with 11 postponements—eventually leaving the retirement eligibility age at 76 so that the current dependency ratio is maintained. Note that even the arguably unpopular policy of raising the retirement age to 76 requires some fiscal adjustments to finance the short-run increases in the dependency ratio.

While nearly all economists agree on the necessity of fiscal adjustments to confront the ageing of the population, only recently have they taken the issue to applied general equilibrium models to inform public discussion of reform proposals. In addition to the fiscal adjustment required to finance social security, there is an additional fiscal burden associated with the ageing of the population. The President's Council of Economic Advisors (1997) projects that medicare and medicaid spending will increase from 2.7 per cent and 1.2 per cent of GDP in 1996 to 8.1 per cent and 4.9 per cent of GDP in 2050, respectively.

This section describes some of the fiscal scenarios considered by De Nardi *et al*. (1999) that the government is likely to follow in dealing with the financing of public pensions and the estimated increases in public health expenditures that are projected to arise with the ageing of the population. The eight fiscal adjustment packages that we summarize have the following features. Experiment (1) places all of the adjustment burden on to the social security payroll tax; experiment (2) raises a consumption tax rate, whereas (3), (4), and (5) reduce benefits in various ways while also adjusting taxes. In particular, (3) postpones the retirement age by 2 years (to reflect the current legislation) and increases the social security payroll tax, whereas (4) postpones retirement in the same way but increases the consumption tax for completing the fiscal adjustment. Experiment (5) explicitly taxes pension benefits while raising the payroll tax to maintain zero social security balance. Experiments (6) and (7) increase the linkage of benefits to cumulative earnings while also adjusting either the labour income or consumption tax rate. Finally, experiment (8) implements a privatization by gradually phasing out benefits while adjusting the labour income tax rate. Except for (8), the experiments leave the social security system unfunded.

Figure 9 shows the welfare cost of implementing policies (2)–(8), relative to policy (1). Our measure of welfare cost is a compensating variation measure which computes a one-time award of assets (as a fraction of assets, measured along the vertical axis) that should be given to people born in a given year (measured along the horizontal axis) to make them indifferent between scenario (1) with compensation and an alternative without compensation. Thus, a positive measure indicates that a positive award would be needed to compensate a person of the indicated birthdate living in scenario (1) to leave them as well off as in an alternative scenario. The figure reveals the different interests served by the different policies. For example, consider an average member of the cohort born in 1940. This individual would rather give up some wealth and stay under the experiment (1) fiscal regime of increasing payroll taxes rather than accept the taxation of benefits of experiment (5).

Overall, Figure 9 summarizes the main findings of De Nardi *et al*. (1999).

- Maintaining pensions at the current levels in the face of projected demographics will require large increases in tax rates. If distorting labour income taxes are used, capital accumulation and labour supply will decline. Back-of-the-envelope accounting calculations made outside a general equilibrium model can be quite misleading and overly optimistic. For example, the Social Security Administration states that a 2.2 percentage point addition to the 12.4 per cent OASDI payroll tax will restore the financial balance in the social security trust fund over the 75-year horizon, given intermediate projections of demographics and other key variables. According to Goss (1998), Deputy Chief Actuary of the Social Security Administration, a 4.7 per cent immediate increase of the existing OASDI payroll tax is necessary to finance the existing social security system in perpetuity.[15]

  When the same population projections are used in our applied general equilibrium model, we obtain an additional 17.1 percentage points in the payroll tax rate and large welfare losses associated with maintaining our current unfunded system. If we include the projected burden stemming from the increase in medicaid and medicare payments, a further 12.7 percentage points will be needed, bringing the total payroll tax including personal income tax to about 60 per cent under policy (1).

- The required fiscal adjustment in the form of higher tax rates is significantly reduced by decreasing the level of pensions through taxation of benefits and consumption or through postponing the retirement eligibility age.

- Although some policies have similar long-run outcomes, there are significant short-run differences in the distribution of generational welfare. For example, gradual privatization leaves most of the current cohorts worse off relative to raising the labour income tax in the future to finance the fiscal burden. Almost all the experiments impose welfare losses on transitional generations, relative to policy (1). The exception is to switch from the current benefit formula, that partially links benefits to contributions, to a defined contribution system. As argued by Kotlikoff *et al*. (1999), fully linking social security contributions to benefits apparently removes a significant distortion from labour–leisure decision.

VII. DIRECTIONS FOR FUTURE RESEARCH

Empirically relevant general equilibrium models of pension and social security reform are computationally intensive, meaning that rapid advances in computing power promise to extend the range of realistic features that we shall be able to incorporate. We briefly mention some important aspects that the next generations of models should treat.

- *Aggregate risk*. Most of the models in the literature assume that there is no aggregate risk. This assumption is made to control the dimension of the state of the economy by keeping cross-section distributions of wealth and income either constant or requiring them to be very smooth functions of time, as in the second experiment described in the previous section. Computationally intensive methods illustrated by Krusell and Smith (1998) and den Haan (1997) raise the promise of incorporating aggregate risk into overlapping generations models like those above. Doing this is important for studying pension reform, because uncertainty surrounding demographic

[15] These optimistic projections do not fully take into account the recent estimates of future government budget surpluses, and therefore can be regarded as upper bounds on the extent of tax increases to confront the ageing of the population.

transitions and rates of return on capital are aggregate risks that can have very important effects on the relative merits of fully funded versus pay-as-you-go social retirement systems.

- *Myopia*. There is a need to digest and import into general equilibrium models of pension reform some of the insights of the work of Strotz (1956), Phelps and Pollak (1968), Laibson (1997), and others on self-control. For one attempt to start this process, see İmrohoroğlu *et al*. (2000), who study the role of social security in an environment with quasi-hyperbolic preferences. In their calibrated general equilibrium model, social security turns out to be a poor substitute for a perfect commitment technology in maintaining old-age consumption. Furthermore, up to a short-term discount rate of 15 per cent, social security is welfare reducing. This is a relatively new area of research and more work is needed to evaluate the quantitative role of institutions related to saving and retirement.

- *Political economy*. We have confined our paper to work that compares the consequences of alternative exogenously specified government policies. Such work can be defended as preparing the way for analyses of political processes that somehow frame and then choose among alternative such policies. Cooley and Soares (1999) and Boldrin and Rustichini (2000) have presented such analyses within the context of some overlapping generations models with democratic voting. Bassetto (2000) performs an analysis in which bargaining among groups, rather than voting, determines taxation and transfer policies.

REFERENCES

Abowd, J. M., and Card, D. (1989), 'On the Covariance Structure of Earnings and Hours Changes', *Econometrica*, **57**(2), 441–55.

Altonji, J. G. (1995), 'Intertemporal Substitution in Labor Supply: Evidence from Micro Data', in O. C. Ashenfelter and K. F. Hallock (eds), *Labor Economics*, *Volume 1*, *Labor Supply and Labor Demand*.

— Hayashi, F., and Kotlikoff, L. J. (1997), 'Parental Altruism and Intervivos Transfers: Theory and Evidence', *Journal of Political Economy*, **105**(6), 1121–66.

Attanasio, O. P. (1994), 'Personal Saving in the United States', in J. M. Poterba (ed.), *International Comparison of Household Saving*, NBER Project Report Series, Chicago, IL, and London, University of Chicago Press, 57–123.

— Hoynes, H. W. (2000), 'Differential Mortality and Wealth Accumulation', *Journal of Human Resources*, **35**(1), 1–29.

— Weber, G. (1995), 'Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey', *Journal of Political Economy*, **103**(6), 1121–57.

— Banks, J., Meghir, C., and Weber, G. (1999), 'Humps and Bumps in Lifetime Consumption', *Journal of Business and Economics Statistics*, **17**(1), 22–35.

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

Avery, R. B., Elliehausen, G. E., Canner, G. B., and Gustafson, T. A. (1984), 'Survey of Consumer Finances', *Federal Reserve Bulletin*, **70**(9), 679–92.

Barro, R. J. (1974), 'Are Government Bonds Net Wealth?', *Journal of Political Economy*, **82**, 1095–117.

Bassetto, M. (2000), 'Political Economy of Taxation in an Overlapping-generations Economy', Federal Reserve Bank of Minneapolis Discussion Paper No. 133.

Boldrin, M., and Rustichini, A. (2000), 'Political Equilibria with Social Security', *Review of Economic Dynamics*, **3**, 41–78.

Browning, M., Hansen, L. P., and Heckman, J. J. (1999), 'Micro Data and General Equilibrium Models', in J. B. Taylor and M. Woodford (eds), *Handbook of Macroeconomics*, Vol. 1a, ch. 8.

Cagetti, M. (2000), 'Wealth Accumulation Over the Life Cycle and Precautionary Savings', University of Virginia, mimeo.

Cass, D. (1965), 'Optimum Growth in an Aggregative Model of Capital Accumulation', *Review of Economic Studies*, **32**, 233–40.

Cooley, T., and Soares, J. (1999), 'A Positive Theory of Social Security Based on Reputation', *Journal of Political Economy*, **107**, 135–60.

De Nardi, M. (1999), 'Wealth Inequality and Intergenerational Links', Federal Reserve Bank of Chicago, Working Paper No. 13.
— İmrohoroğlu, S., and Sargent, T. J. (1999), 'Projected US Demographics and Social Security', *Review of Economic Dynamics*, **2**(3), 575–615.
Den Haan, W. J. (1997), 'Solving Dynamic Models with Aggregate Shocks and Heterogeneous Agents', *Macroeconomic Dynamics*, **1**(2), 355–86.
Diamond, P. A. (1977), 'A Framework for Social Security Analysis', *Journal of Public Economics*, **8**, 275–98.
— Hausman, J. A. (1984), 'Individual Retirement and Savings Behavior', *Journal of Public Economics*, **23**(1/2), 81–114.
Dynan, K. E., Skinner, J., and Zeldes, S. (2000), 'Do the Rich Save More?', NBER Working Paper No. 7906, September.
Fernández-Villaverde, J., and Krueger, D. (2000), 'Consumption and Saving over the Life Cycle: How Important are Consumer Durables?', Federal Reserve Bank of Minneapolis, manuscript, October.
French, E. (2000), 'The Effects of Health, Wealth and Wages on Labor Supply and Retirement Behavior', Federal Reserve Bank of Chicago Working Paper 2000-02.
Fuster, L. (1999), 'Is Altruism Important for Understanding the Long-run Effects of Social Security?', *Review of Economic Dynamics*, **2**(3), 616–37.
Goss, S. C. (1998), 'Measuring Solvency in the Social Security System', manuscript presented at the 1997 Pension Research Council Symposium at the Wharton School of the University of Pennsylvania, Baltimore, MD, Social Security Administration.
Gourinchas, P.-O., and Parker, J. A. (1999), 'Consumption over the Life Cycle', NBER Working Paper No. 7271.
Gruber, J., and Wise, D. (1999), *Social Security Programs and Retirement Around the World*, NBER Conference Report Series, Chicago, IL, University of Chicago Press.
Hansen, G. D. (1993), 'The Cyclical and Secular Behavior of Labor Input: Comparing Efficiency Units and Hours Worked', *Journal of Applied Econometrics*, **8**(1), 71–80.
Hayashi, F. (1995), 'Is the Japanese Extended Economy Altruistically Linked? A Test Based on Engel Curves', *Journal of Political Economy*, **103**(3), 661–74.
Huang, H., İmrohoroğlu, S., and Sargent, T. J. (1997), 'Two Experiments to Fund Social Security', *Macroeconomic Dynamics*, **1**(1), 7–44.
Hubbard, G. R., and Judd, K. L. (1987), 'Social Security and Individual Welfare', *American Economic Review*, **77**(4), 630–46.
— Skinner, J., and Zeldes, S. P. (1995), 'Precautionary Saving and Social Insurance', *Journal of Political Economy*, **103**(2), 360–99.
Huggett, M., and Ventura, G. (1999), 'On the Distributional Effects of Social Security Reform', *Review of Economic Dynamics*, **2**(3), 498–531.
Hurd, M. D. (1987), 'Savings of the Elderly and Desired Bequests', *American Economic Review*, **77**(3), 298–312.
İmrohoroğlu, A., İmrohoroğlu, S., and Joines, D. H. (1995), 'A Life Cycle Analysis of Social Security', *Economic Theory*, **6**, 83–114.
— — — (1999), 'Social Security in an Overlapping Generations Economy with Land', *Review of Economic Dynamics*, **2**(3), 638–65.
— — — (2000), 'Time Inconsistent Preferences and Social Security', Federal Reserve Bank of Minneapolis Discussion Paper 136, August.
Koopmans, T. C. (1960), 'Stationary Ordinal Utility and Impatience', *Econometrica*, **28**, 287–309.
Kotlikoff, L. J., Smetters, K., and Walliser, J. (1999), 'Privatizing Social Security in the United States—Comparing the Options', *Review of Economic Dynamics*, **2**(3), 532–74.
Krueger, D., and Perri, F. (2000), 'Does Income Inequality Lead to Consumption Inequality? Empirical Findings and a Theoretical Explanation', Stanford University, manuscript.
Krusell, P., and Smith, A. (1998), 'Income and Wealth Heterogeneity in the Macroeconomy', *Journal of Political Economy*, **106**, 867–96.
Kydland, F. E., and Prescott, E. C. (1996), 'The Computational Experiment: An Econometric Tool', *Journal of Economic Perspectives*, **10**(1), 69–85.
Laibson, D. (1997), 'Golden Eggs and Hyperbolic Discounting', *Quarterly Journal of Economics*, **112**, 443–77.
McGrattan, E., and Rogerson, R. (1998), 'Changes in Hours Worked Since 1950', *Federal Reserve Bank of Minneapolis Quarterly Review*, **22**(1), 2–19.
MaCurdy, T. E. (1982), 'The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis', *Journal of Econometrics*, **18**(1), 83–114.

Palumbo, M. G. (1999), 'Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle', *Review of Economic Studies*, **66**(2), 395–421.
Phelps, E. S., and Pollak, R. A. (1968), 'On Second-best National Saving and Game-equilibrium Growth', *Review of Economic Studies*, **35**, 185–99.
Ramsey, F. P. (1928), 'A Mathematical Theory of Saving', *The Economic Journal*, **38**, 543–59.
Samuelson, P. A. (1958), 'An Exact Consumption–Loan Model of Interest with or without the Social Contrivance of Money', *Journal of Political Economy*, **66**(6), 467–82.
Shorrocks, A. F. (1975), 'The Age–Wealth Relationship: A Cross-section and Cohort Analysis', *Review of Economics and Statistics*, **57**(2), 155–63.
Solow, R. M. (1956), 'A Contribution to the Theory of Economic Growth', *Quarterly Journal of Economics*, **70**, 65–94.
Storesletten, K., Telmer, C. I., and Yaron, A. (2000), 'Consumption and Risk Sharing over the Life-Cycle', mimeo, http://www.iies.su.se/~storeslk/
Strotz, R. H. (1956), 'Myopia and Inconsistency in Dynamic Utility Maximization', *Review of Economic Studies*, **23**, 165–80.
Velde, F. (1999), 'Americans Are Not Saving: Should We Worry?', *Chicago Fed Letter*, No. 141, May.
Villanueva, E. (1999), 'Parental Altruism under Imperfect Information: Theory and Evidence', Barcelona, Pompeu Fabra, mimeo.