forecasting and performing analysis with the model is to provide detailed information to decision-makers based on the model's rich structure. Their procedure is to incorporate projections of macroeconomic variables, treating these as exogenous, and then to use the model structure to generate consistent forecasts for specific sectors. The Australian motor vehicle industry is the case study presented. Given the forecast baseline, the effect of a change in a policy instrument may be analyzed using a standard simulation procedure. This procedure implies that the effect of a given policy will typically depend on the macroeconomic forecast that the model is calibrated to reproduce. As an illustration, the effect of a reduction in the tariff on imported cars from 12% to 4% is considered. This policy reduces domestic vehicle production and increases imports of vehicles. Aggregate macroeconomic effects from the policy are almost completely absent, and the changes in the production structure are such that losers are mainly found in industries with a heavy dependence on the vehicle industry.

References

Auerbach, Alan J. and Kotlikoff, Laurence J., *Dynamic Fiscal Policy* (Cambridge, UK: Cambridge University Press, 1987).

Using Dynamic General Equilibrium Models for Policy Analysis – G.W. Harrison et al. (Editors)
© 2000 Elsevier Science B.V. All rights reserved.

2. The A-K OLG Model: Its Past, Present, and Future

by
Laurence J. Kotlikoff

1. Introduction

It is now almost two decades since the development of the prototype of the Auerbach-Kotlikoff dynamic life-cycle simulation model. The model has been used to examine a host of policies, including tax reform, tax cuts, investment incentives, tax progressivity, expansion of social security, government spending, monetary policy, endogenous growth, the size of the informal sector, human capital accumulation, and educational policy. It has also been used to study demographic change, the timing of policy impacts, the efficiency gains from fiscal reforms, and the effects of fiscal policies on both the intra- and intergenerational distributions of economic welfare.

Although much of the research on the model has been conducted by Auerbach and Kotlikoff and their co-authors, economists around the world in academia, government, and multilateral lending institutions have developed and studied their own versions of the model.

This paper describes, in the next section, the origins of the A-K Overlapping Generations (OLG) Model and the economic profession's reaction to it. Section III considers some of the general lessons that have been learned in using the model. Section IV illustrates the model's current capacities by considering tax reform and social security privatization. Section V discusses outstanding research questions and the need for additional improvements to the model. The Appendix, reproduced from Kotlikoff, et al. [1997], lays out the model's current structure and describes its method of solution.

2. Origins of the A-K OLG Model

The 1970s was a period of tremendous intellectual ferment in academic economics, in general, and in public finance, in particular. In macroeconomics, Lucas and Sargent were dismantling the foundations of postwar Keynesian economic analysis. Their critique focused on shortcomings in Keynesian econometric practice. But in forming their critique, they stressed the importance of connecting macroeconomic outcomes to microeconomic fundamentals and that the current economic decisions of rational economic agents depend as much on their expectations about the future as they do on their immediate circumstances.

The Rational Expectations Revolution made the entire profession think more about dynamics, particularly about the dynamics of neoclassical (micro-based) models. The foundation for considering models of economic change over time had, of course, been laid in previous decades by Domar, Harrod, Koopmans, Samuelson, Solow, Diamond, Phelps, Stiglitz, Cass, Shell, Sidrauski, Tobin, Uzawa and numerous other growth theorists. So most of the theoretical machinery was well in place by the time Diamond, Feldstein, Stiglitz, Sheshinski and other public finance economists started addressing the issue of the dynamic impacts of various fiscal policies.

Feldstein's contribution here in stimulating research on fiscal policy and growth cannot be overstated. His debate with Robert Barro [1974] about the impact of government debt and unfunded social security fascinated a generation of graduate students who realized that fiscal policy could be a lot more interesting that simply shifting an IS curve. Feldstein [1974] not only rekindled a two century-long debate about the burden of the debt, he also showed that explicit government debt policies were not unique in their ability to generate large intergenerational transfers. Finally, in stressing social security's incentives for early retirement, Feldstein raised the more general question of how the entire panoply of effective marginal taxes and transfers could affect an economy's long-run health. Another key contribution by Feldstein [1977] during this period was his study of the burden of a land rent tax. This paper demonstrated that, by altering the market valuation of capital assets, fiscal policy can could effect intergenerational redistributions and alter the economy's growth path and long-run rest point.

Feldstein's theoretical work was couched in simple models and provided a qualitative sense of how fiscal policies could influence the economy. Although his empirical work provided some quantitative feel for the potential impacts of particular policies, Feldstein left five major issues unresolved.

First, what were the long-run quantitative effects of fiscal reforms in more realistic multiperiod models? Second, how long did it take to get to the long-run? Third, were short run policy effects necessarily of the same sign as long run effects? Fourth, what was the structural relationship between alternative fiscal policies that effected quite similar intergenerational redistributions, and why did they eventuate in such different measures of official government debt? Fifth, how much of the potential long-run improvement in an economy's well being arising from fiscal reforms was the result of efficiency improvements as opposed to the sacrifice of those alive in the short and intermediate runs?

The revolutions underway in macroeconomics and public finance were concomitant with the revolution in the access of academic economists to high speed computers. Although PCs (personal computers) did not yet exist, computer programs could be punched onto computer cards and run on university mainframes. John Shoven and John Whalley [1972] were perhaps the first public finance economists to demonstrate the potential power of this new technology to study economies in general equilibrium. They showed how Harberger's model of the corporate income tax could be expanded to include multiple industries, each facing its own effective rate of corporate income taxation, and solved, using Scarf's algorithm, on the computer.

Shoven-Whalley Model was static, but it demonstrated that one could use the computer to find market-clearing prices of multiple interconnected markets that were open at the same point in time. Although solving for the economy's dynamic transition path involved finding a general equilibrium with respect to markets open at different points in time, this distinction didn't necessarily seem important. After all, Arrow's contingent claims model showed that markets that are indexed by time can effectively be considered open in the present. On the other hand, in Arrow's model all the agents who would ever be alive were already alive and were, therefore, available to form contracts in the present about what they would do in the future. This is not the case in an ongoing life-cycle economy in which current agents are not altruistic toward future ones and will not, as a result, negotiate on their behalf. However, even this concern seemed surmountable. After all, one could consider how the agents to be born in the future would contract once they arrived. And given the capacity of computers, there seemed, in principal, to be no problem in simultaneously entertaining the economic behavior of those now alive as well as those yet to be born. By limiting oneself to economies experiencing no aggregate uncertainty, one could hope

to narrow all future paths of the macro economy to just one and therefore not have to entertain actions by future agents that differed depending on the state of the economy.

The real issue in immediately applying the Shoven-Whalley procedure to markets over time was that the Scarf algorithm only worked for a model with a finite number of markets and agents. There was no obvious way to solve for a general equilibrium with an infinite number of time-dated agents transacting in an infinite number of labor, consumption, and capital markets. This is a bit of an overstatement, since in certain very simple two-period life-cycle models (in which households live for just two periods), solving for the exact transition path is simple because next period's capital-labor ratio is simply a function of this period's capital-labor ratio — a known quantity. Once one adds an extra period or makes the two period model a bit more complicated, the capital-labor ratio in any given period depends on the capital-labor ratio in at least the previous two periods. Stated differently, knowing this period's capital-labor ratio doesn't suffice to determine next period's ratio. One also needs to know the capital-labor ratio that will arise in future periods. But the equations containing the capital-labor ratios in those future periods involve capital-labor ratios in periods that lie yet further in the future. Hence, substituting from these equations for the unknown ratios just adds different unknown ratios to the resulting expression.

Mathematically, a three-period life-cycle model gives rise to a third- or higher order nonlinear difference equation for which there are, in general, no known analytical solutions. The problem becomes even more complicated in dealing with 55 generations. In the A-K OLG Model agents consider factor prices and, thus, capital-labor ratios, over their entire lifetimes. Hence, the youngest agent at any point in time is considering, among other things, the factor prices and capital-labor ratio that will prevail 55 years in the future. But this agent knows that the capital-labor ratio 55 years in the future will be determined, in part, by the labor supply of the youngest agent in the economy in that year. This youngest agent 55 years hence will, in turn, be considering factor prices and, thus, capital-labor ratios, over the following 55 years. In sum, the economic decisions of the youngest agent alive this year depend on 110 capital-labor ratios; i.e., the A-K OLG Model entails, roughly speaking, a 110th order nonlinear difference equation for the economy's key state variable — it's ratio of capital to labor.

Now one could, as Laitner [1984] ultimately did, linearize such a high order non linear difference equation and apply standard solution techniques

from the theory of linear difference equations. But in so doing, one would be studying an approximation to the actual transition path of the economy under investigation. Without solving for the exact transition path, one wouldn't be able to tell the precise size of the approximation error.

Absent a method of solving for the economy's exact transition path, public finance economists in the late 1970s turned to solving for economies' long-run steady-states. The steady states of dynamic economies also entail an infinite set of agents and markets — but the same set over time. This feature provides the mathematical closure for computing where certain large-scale economies would end up even if one didn't know precisely how they would get there. Two examples of the line of research are Kotlikoff [1979][1], which considered the long-run effects of pay-as-you-go social security, and Summers [1981][2], which examined the long-run effects of tax reform. Both of these studies considered 55-period life cycle models in which adult agents live from age 20 through age 75. Kotlikoff's study confirmed Feldstein's contention that unfunded social security, of the scale established in the U.S., would have a significant deleterious impact on an economy's ultimate capital stock and living standard. Summers's study showed that the choice of the tax base used to finance government spending could have equally large long-run impacts on the economy.

Knowing that the fiscal policies matter greatly in the long run raised the ante in studying the transition, particularly the duration of the transition. If the long-run took forever to reach, then knowing the economy's ultimate position wasn't particularly useful. On the other hand, if the transition to the economy's final destination was very quick, then steady-state analysis would suffice. Either way, one would need to compute transition paths.

Summers [1981] and Seidman [1983], following the lead of Miller and Upton [1974], took a stab at this by computing transition paths based on the assumption of myopic expectations, specifically the assumption that agents assume that current factor prices will prevail at all future dates. This approach seems appealing on first thought, but not on second or third thought. First, one has to ask why agents would consider simply current, as opposed to past, factor prices in thinking about future factor prices. Second, if one permits agents to think about past factor prices, how many past prices should one let agents consider? Third, what should one assume about the way agents weigh current and past factor prices in forming expectations about future factor prices? Fourth, should one value agent's well being based on the actual time-path of factor prices that prevails or the one that agents mistakenly assume will prevail? And fifth, how do agents consider

fiscal actions taken today, such as a tax cut, that necessitate future fiscal actions, like a spending cut, to satisfy the government's intertemporal budget constraint? Ignoring general equilibrium feedback effects on factor prices is one thing, but assuming that agents think the government can borrow indefinitely to pay its bills is something else again.

In 1980 Martin Feldstein announced a major NBER conference on simulation methods in tax policy analysis to be held in 1981. Not knowing what we would write, but knowing that the conference would be held at a fancy hotel in Florida in January, Auerbach and I agreed to submit a paper. In the winter of 1980 we met to figure out what to do. We began talking about the need to compute exact transition paths in large-scale life-cycle economies under the assumption of perfect foresight and the fact that we had no idea how to do so. Then we starting thinking about the method for solving for steady states. Since the equation for the economy's steady-state capital-labor ratio isn't a closed form, solving for its value involves iteration — trying a value, testing if it solves the steady state restrictions, and updating one's guess if it doesn't. More precisely, the procedure involves choosing a candidate value for the steady-state capital-labor ratio demanded by firms, using the marginal product equals factor price equations to determine factor prices, and then calculating whether the supplies of capital and labor provided by the household sector of the economy, given these factor prices, produce the same ratio of capital to labor as the candidate demand value. If not, one updates the candidate value by taking a weighted average of the initial demand and the new supply-side values. An alternative to this Gauss-Seidel technique is to use Newton's method to update one's guess of the solution. After contemplating for a couple of days the embarrassment of showing up at the conference with no paper, it suddenly struck us that this same procedure used to solve for steady states might work in solving for the economy's exact transition path. In the case of solving for the transition path, one would guess the time-path of the amount of capital relative to labor demanded by firms, use this path to calculate a time-path of factor payments, and then use this time-path of factor payments to determine a time-path of the supply of capital relative to labor. If one found a time-path of demands of capital relative to labor that equaled the time-path of the supply of these factors, one would have calculated a dynamic equilibrium.

At this point, we had two questions left to worry about: How should we handle the fact that our economy had no terminal period? And, would our procedure of solving for the transition path converge? To deal with the terminal period issue, we decided to assume that the economy reached a steady state at a date that was sufficiently far off in the future that the economy would, indeed, have plenty of time to reach that steady state. Thus, in forming our initial guess of the time-path of the demand for capital relative to labor and in updating this guess, we assumed that the value of this ratio was constant from year 150 onward; i.e., we guessed the same value of this ratio for years after 150 as for year 150.

To deal with the second question — Would our procedure converge? — we simply had to try it. In those days (the winter of 1980), computer technology had improved, so one no longer had to punch out computer cards, but could instead use a terminal that looked like a large typewriter and had no screen. Your results came back by way of the terminal typing them out on paper. Even better, one could run this from home via a big modem into which one stuck the terminal's phone. Alan and I spent a couple of days writing a prototype program in my Los Angeles apartment. We hadn't debugged it by the time Alan had to head back East. So for the next couple of days I sat in front of this machine trying to remove the final bugs and see if we could get the right sequence of numbers to come out. What we were printing out was our successive time-paths of the capital-labor ratio and the difference between the values of the current time-path and the previous one. To me it seemed that we were looking at time-lapsed photographs of a wiggly snake. Each successive picture of the snake (of the time-path of the capital-labor ratio) showed different wiggles, with the question being would the successive snakes start flattening out — start lining up on top of each other. After fixing what I thought had to be our final programming bug, I sat back to watch the snake and sure enough its position started to flatten out and, what's more, become more flat after each iteration. When the snake had flattened at all points (at all time periods) to a degree that was well within our convergence criteria and I saw that convergence was improving across the snake's entire length in each iteration, I removed the phone from the modem and called Alan. He wasn't home, but I left a message with his wife (there were no answering machines back then) that "The snake has flattened." This was an historic moment for us because we realized two things: first, we'd have a paper for the conference and second, we'd probably get tenure.

The first version of the A-K OLG Model had exogenous labor supply and proportional taxes. In the period right after the conference we teamed up with Jon Skinner to add variable labor supply, endogenous retirement, progressive taxes, and a lump-sum redistribution authority that could help us distinguish efficiency gains from intergenerational redistribution. This

collaboration resulted in Auerbach, Kotlikoff, and Skinner [1983]. In the mid1980s, Alan and I added investment incentives and quadratic costs of adjusting the capital stock. As a result, we had a model that we could use to see how intergenerational redistributions could be achieved via policy-induced asset market revaluations and how the stock market's value would adjust in response to policy changes. We also added a social security system and, in a simplified version of the model, demographics, including the presence of children whose consumption entered in their parents' utility functions when they were children. All of these features were included in our book *Dynamic Fiscal Policy* that we published in 1987 with Cambridge University Press.

After teaming up with Robert Hagemann and Guiseppe Nicoletti (who were at the OECD at the time) to study demographic change in four OECD countries (Auerbach, Hagemann, Kotlikoff, and Nicoletti [1989]), Alan and I took a break from dynamic fiscal simulations in the early 1990s to work on generational accounting and other issues. But we always knew that the simulation model had lots more potential. I started using the model again a couple of years ago to study the degree to which generational accounting was approximating true intergenerational fiscal incidence (Fehr and Kotlikoff [1996]) and the privatization of social security (Kotlikoff [1996]). At that point it became clear that the biggest shortcoming in the model was its failure to address intragenerational equity issues. But it also seemed that Fullerton and Rogers [1993] had already addressed that issue. Their book, which contained simulations based on a dynamic life-cycle simulation model, seemed to deal with both intragenerational and intergenerational issues. The Fullerton-Rogers model also featured lots of goods (our model had just one) and lots of industries producing these goods.

In considering whether we'd be simply reinventing Fullerton and Rogers, I realized that, for all its attractive features, their model lacked one essential element -- the ability to solve for the economy's perfect foresight transition path; i.e., the Fullerton-Rogers model used the same myopic expectation assumption adopted by Summers back in 1980.

So I hooked up with Kent Smetters and Jan Walliser, who were both at the Congressional Budget Office, to add intragenerational heterogeneity to the model. We decided to follow very closely Fullerton and Rogers' innovative approach of specifying 12 different human capital ability groups within each cohort. Kent and Jan did the yeoman's job of reprogramming the A-K OLG Model, after which we started writing papers on social security

reform (Kotlikoff, Smetters, and Walliser [1997, 1998a, 1998b, 1998c]). During this period, I was also talking with David Altig who, together with Chuck Carlstrom, had done simulation work on a model that featured nondifferentiable budget constraints where the nondifferentiabilty arose from the discrete brackets of the federal income tax. I suggested that David team up with Alan, me, Kent, and Jan to work on a simulation study of U.S. tax reform that included non differentiable budget constraints as well as intragenerational heterogeneity. This collaboration resulted in Altig, et al. [1997].

Before pointing out the key lessons we've learned over the years with the A-K OLG Model, illustrating the model's current capacities, and outlining further improvements needed in the model's structure, let me say a couple of words about the model's acceptance by the economics profession.

The Reaction to the A-K OLG Model

Reaction to the A-K OLG Model has generally been quite favorable. Stiglitz, who discussed on our first paper at the 1981 simulation conference, referred to it as "a tour de force." And Lucas [1990] gave the work a strong endorsement. But other economists voiced concerns. Some felt that simulating policies was intrinsically different from using the calculus to sign and size their effects. Some referred to the model as a "black box," whose results could neither be verified nor fully understood. And some (e.g., Kehoe and Levine [1985]) questioned whether our model had a unique equilibrium (a unique transition path) and suggested that our fixing of a year by which the economy reaches a steady state might be forcing the model to pick one out of a potentially very large number of transition paths.

I've always viewed the concerns as misplaced. At a conceptual level, it's hard to understand how finding an exact solution to a set of equations and, therefore, being able to determine the precise response to potentially large policy changes could be worse than using the calculus. Taking and evaluating derivatives is only valid for very small policy changes. Furthermore, the evaluation of derivatives of any economic variable at any point in time with respect to some policy change requires, in general, knowing the initial (pre-policy change) values of all the economic variables at that point in time. If the economy is in a steady state initially, one can use steady-state values for evaluating such derivatives because the initial (pre-policy change) position

of the economy at any point in time is the steady state position. But what if the economy is moving along a transition path? Then the derivatives of economic variables at any particular point in time must be evaluated using the values of the economy's initial (pre-policy change) variables at that point in time. But how does one know these values unless one actually computes the economy's dynamic equilibrium?

In general then, one can not use the calculus to study the impact of policy changes on the economy through time unless one computes the initial (pre-policy-change) transition path. But if one can compute the initial transition path (the transition path under the initial policy), one can also compute the new transition path (the transition path under the new policy) and compare them. So the calculus is really of no use. Furthermore, restricting oneself to evaluating derivatives for economies that are initially in a steady state is only valid for steady states that can be computed independently of the transition path. But some steady states can not be so computed, specifically steady states in which the value of long-run policy variables are set as functions of what happens during the transition. As an example, take an economy in a steady state with no social security. Now suppose the government announces that it is going to establish a pay-as-you-go social security system starting in five years by setting, on a permanent basis, social security's benefit at 40% of the average wage prevailing in five years. In order to figure out what that benefit level and its associate payroll tax are, one has to calculate the transition and, among other things, determine the level of average wages in year 5.

The complaint about the model being a "black box" is puzzling given that there are only a limited number of elements in the model — a standard, time-separable CES utility function, a standard CES production function, standard quadratic capital adjustment costs, a familiar fiscal structure, a few other quite straightforward elements, and no uncertainty. If this most basic of neoclassical life-cycle growth models, constructed along a realistic time dimension, is viewed as a "black box," it's certainly high time to study that box. The "black box" pejorative is also inappropriate given that the results of the policy simulations are readily explained in terms of economic theory, invariably coincide with one's economic intuition, and are easy to double check.

Although Alan and I have never formally described our methods of double checking the code, after each modification of the model we subject it to a set of tests to make sure there are no misspecifications (bugs). One of these tests is checking that national saving is precisely zero in the long-run if

there is neither population nor productivity growth. Another is to check that the economy sits in its initial steady state when one a) runs a transition, but b) specifies no policy change. A third test is to check that each agent exactly exhausts her budget constraint at the end of her life and never violates her non negativity constraint with respect to labor supply. A fourth is to check that the government is satisfying its intertemporal budget constraint. And a fifth test is to confirm that alternative ways of running the exact same fiscal policy, such as taxing consumption via a retail sales tax or via a proportional income tax with 100% expensing, produce precisely the same economic results. Many of these checks are redundant; i.e., if any agent was off her budget constraint or if the government was off its budget constraint, the economy's steady-state saving rate would not be zero when zero rates of population and productivity growth are assumed.

The final concern — that our economy might feature multiple equilibria — was something Alan and I never worried about. The reason is that after doing hundreds of simulations with a range of different initial conditions, we knew instinctively that, for the parameter values we were using, the model was finding unique equilibria. Why were we so sure? The answer is that if the model featured multiple equilibria, the transition path would be a) either very hard to compute, b) highly sensitive to initial conditions, or c) highly sensitive to the year after which we assumed the economy was in a steady state. Again, for our assumed preference and technology parameters we didn't find any of these things to be the case. Moreover, when we set particular parameter values to the extreme levels that foster multiple equilibria, such as very small values for the intertemporal elasticity of substitution or the elasticity of substitution in production between capital and labor, we found we were no longer able to compute transition paths or even steady states.

If our intuition and simulation experience told us that our transition paths were unique, we were still very relieved when Laitner [1984] took on the formidable task of examining this issue formally by linearizing our model. He found that, for at least the linearized version of our model, there was a unique transition path for the range of parameter values we were using.

Of the various reactions to our model, the one that has been most gratifying is its apparent influence on the modeling strategies of other researchers. Hamann [1992], Arrau and Schmidt-Hebbel [1993], Raffelhuschen [1989, 1993], He Huang, Selo Imrohoroglu, and Thomas Sargent [1997], Imrohoroglu, Imrohoroglu, and Joines [1995, 1998a, 1998b], Altig and Carlstrom [1996], Heckman, Lochner, and Taber [1998a,

1998b], Hirte and Weber [1998], Schneider [1997], Fougere and Merette [1998, 2000], Merette [1998], Lau [2000], Rutherford [2000], and Schmidt-Hebbel [1999] are examples in this regard. Equally gratifying, the model is starting to get attention in Washington. The Congressional Budget Office is now using the model to study both tax reform and privatizing social security, and the Joint Committee on Taxation of the U.S. Congress has included the model's results in its recent analysis of the economic gains from tax reform (see the Joint Committee on Taxation [1997]).

Our presentation of and emphasis on the government's intertemporal budget constraint has helped clarify the fact that, apart from potential efficiency gains, the welfare effects of policy changes are zero sum in nature — the losses to the losers (including those coming in the future) have the same value in the present as the gains to the winners (including those coming in the future). Indeed, Alan, Jagadeesh Gokhale and I have made the government's intertemporal budget constraint the central structure of generational accounting. Alan and my presentation of the government's intertemporal budget constraint in our original 1980 working paper was probably not the first such presentation, but it was surely one of the earliest.

Finally, I think the model has made a real contribution in distinguishing welfare changes arising from efficiency improvements from those arising from intergenerational redistribution. Too many public finance economists have substituted their own value judgements for economic science in measuring the net welfare gains from policy changes. The all-too-frequent practice is simply to choose one's own discount rate in discounting policy-induced welfare changes experienced by current and future generations and calling the net total a "welfare gain." In so doing, one is imposing one's own social welfare function (one's own ethical norms) without acknowledging that to be the case; i.e. there is seldom any clarification in these studies that their authors are using their own ethical judgements in making welfare statements and welfare comparisons.

3. Lessons Learned

What have we learned from the A-K OLG Model? The answer is some things that we expected and some things we did not. As expected, we confirmed that the choice of the tax structure, the progressivity of income tax rates, the expansion of pay-as-you-go social security benefits, and deficit-financed tax cuts can very substantially alter an economy's long-run well being. We also confirmed, as Chamley's [1981] seminal work had suggested would be the case, that consumption taxation was more efficient than either income taxation or wage taxation and that capital income taxation by itself was the less efficient than either income taxation, consumption taxation, or labor income taxation.

What we didn't expect to find is that economic transitions are really quite slow. Policies that, say, reduce the economy's capital stock by a third in the long run can have a half life of 20 to 30 years. We were also surprised to find that the duration of a policy matters a great deal to its short-run impact. Take deficit-financed tax cuts. If these are of short duration, substitution effects dominate income effects and the economy expands as workers take advantage of temporarily low tax rates. The short-run expansion can be quite substantial, not so much because of the size of compensated substitution effects (our parameter values are fairly conservative), but because of the magnitude of the size of the changes in economic incentives. So the supply side zealots in the U.S. are partly right; deficit-financed tax cuts can raise saving, crowd capital, and expand output. What they get wrong is that these gains are short lived. Although short-term tax cuts lead to short-run crowding in, they lead to long-run crowding out and a lower long-run level of capital and per capita output than initially prevailed. Why? Because the economy's tax base doesn't expand in the short run by enough to permit a permanent reduction in tax rates. When the end of the period of the tax cut is reached, the government must raise tax rates, above their pre-policy level, in order to cover its spending plus interest payments on the new debt it incurred up to that point in time. This was true in our simulations no matter what tax rates were being cut, including capital income tax rates.

Related to our findings about deficits was our general finding that income (redistribution) effects matter much more than substitution (incentive) effects in determining the saving and growth effects of fiscal policies. Basically, if you cut someone's taxes, he's going to spend that tax cut. He may work more and save more while the tax rates are low, but he is all the time planning to spend the tax break once the rates go back up. Moreover,

his short-term reduction in consumption while the rates are low will be less than would be the case had he not experienced the increase in lifetime net income. So during the tax cut period, income effects are limiting the economy's saving response and precluding the tax base from rising by enough to obviate having to raise tax rates again in the future.

The income effects are particularly strong for the initial elderly. Cutting their net taxes by either lowering their gross taxes or providing them with more benefits constitutes an intergenerational redistribution to them because when the government gets around to paying for those net tax cuts, the elderly will either be deceased or, if alive, have relatively few years left to live. Either way, the elderly avoid the tax increase. In the life cycle model, the elderly have larger propensities to consume than do the young because they have fewer years left to live over which they need to spread any increase in remaining lifetime income.[3] Those not yet born have the smallest propensity to consume – zero, because they're aren't yet here. So redistributing to the current elderly from current young and future generations is, in effect, transferring income from low spenders (savers) to high spenders (dissavers).

In addition to teaching us that transitions are slow and that short-run effects can have the opposite sign as long-run effects, the simulations showed us the structural similarities between seemingly disparate policies. Take a policy of making the income tax more progressive in a revenue neutral manner. Although one's first instinct is to view this policy in terms of its economic disincentives, much of the reason that the policy harms saving and short-term economic growth is that it shifts the tax burden away from elderly spenders, who have relatively low current income, and onto younger savers who have relatively high current income because they are still in the work force. So changes in the degree of tax progressivity affect the economy in large part through intergenerational redistribution. In this respect, the policy has a lot in common with deficit-financed tax cuts, pay-as-you-go social security, and changes in the tax base from, for example, consumption to wage taxation.

Intergenerational redistribution also turns out to play a critical role in the expansion of investment incentives. Such a policy leads to a decline in the market value of existing capital, which, because they've had more time to accumulate wealth, is owned primarily by the elderly. The reason is that existing capital doesn't qualify for investment incentives, but must, nonetheless, compete in the marketplace with new capital. The capital loss suffered by the elderly as a consequence of rasing investment incentives is

matched by an effective capital gain to the young and future generations who now are able to purchase the existing capital stock at a lower price. One question that we had when we saw these results is whether they would be reversed if we added capital adjustment costs. The answer turned out to be no, although capital adjustment costs did reduce by about one third the magnitude of the initial decline in the stock market revaluation arising from investment incentives.

The common structure (the common income effects) of seemingly disparate policies kept me wondering through the early 1980s about our fiscal nomenclature. How is it that we describe such similar policies so differently? And why do some policies that hurt young and future generations show up as budget deficits, whereas others do not? Feldstein [1974] had planted the seed for worrying about fiscal taxonomy in the 1970s in calculating unfunded social security liabilites and, thereby, implicitly comparing them with explicit liabilities. But he never clarified why it was that some liabilities that we leave for our children show up on the books, while others do not.

One night, while lying in bed on vacation, it dawned on me that the different policies we had simulated were not really different, rather we were simply using different words to describe the same policies. For example, in the case of pay-as-you-go social security we call contributions being made to the government "taxes" and the benefits received "transfer payments." But we could just as well call the contributions "loans from us to the government" and the benefits "return of principal plus interest on those loans less an old-age tax." In the case of the U.S., this alternative language would roughly triple the size of official government debt. I considered waking up my wife to tell her this, but thought better of it. Instead, I wrote a series of papers on deficit delusion, and Alan and I devoted a chapter of our book on the A-K OLG Model to the issue of deficit finance and fiscal illusion. The most important of these papers is Kotlikoff [1993] which points out that in any neoclassical model with rational economic agents and rational economic institutions (institutions that produce the same economic environment regardless of the language used to describe that environment), the "deficit" is not a well defined economic concept. Stated differently, when one uses " the" deficit to describe a country's fiscal policy, one is, in fact, describing its vocabulary, rather than anything about its actual policy.

A final lesson we learned was the importance of announcement effects to economic efficiency and short-run economic performance. When we orginally wrote our model we didn't realize that it was capable of analyzing

policies that are announced in the present to take place in the future. But it was indeed. All we needed to do was to tell the model when the policy was going to take place, and the agents in the model would automatically and fully incorporate this information in their economic plans. Thus an anouncement today that the economy is switching to a consumption tax two years hence stimulates lots of consumption right now as agents seek to purchase consumption while its still relatively cheap. The short-run run effect on consumption produces the opposite effect of what consumption-tax advocates seek and advertise — an immediate increase in national saving and capital formation. Moreover, the distortion of short-run relative prices can be so great as to either substantially offset or fully eliminate the efficiency gains from consumption taxation.

4. Illustrating the A-K OLG Model

Tax Reform

The U.S. is engaged in an ongoing debate about whether and how to transform its hybrid system of income and consumption taxation into a clean, simple tax system. A variety of alternatives to the current system have been proposed including switching to a clean (no deductions) proportional income tax, a clean proportional consumption tax, a flat tax (that taxes business cash flow plus wage income above a certain level), a progressive personal consumption tax, and David Bradford's X Tax (that taxes wages at progressive rates and business cash flow at the top rate applied to wages).

Altig, Auerbach, Kotlikof, Smetters, and Walliser [1997] found that switching from the current U.S. tax system to a clean proportional consumption tax raises the long-run capital stock by almost one third and long-run output by 11%. Future middle- and upper-income classes gain from this policy, but inital older generations are hurt by the policy's implicit wealth tax. Poor members of current and future generations also lose. The flat tax, in contrast, improves the well-being of all members of future

generations, but at a cost of halving the economy's long-run output gain and harming initial older generations. Insulating these older generations through transition relief further reduces the long-run gains from tax reform. Switching to a proportional income tax without deductions and exemptions hurts current and future low lifetime earners, but helps everyone else. It also raises long-run output by over 5%. Finally, the X tax makes everyone better off in the long run and also raises long-run output by 8%.

Figure 1 and Table 1, reproduced from Altig, et al. [1997], illustrate the model's predictions for the case of switching to a clean consumption tax. Note that income class 1 refers to the poorest members of each cohort (those with the smallest endowment of human capital) and income class 12 refers to the richest members of each cohort (those with the largest endowment of human capital).

Figure 1: Proportional Consumption-Tax Reform, Remaining Lifetime Utility

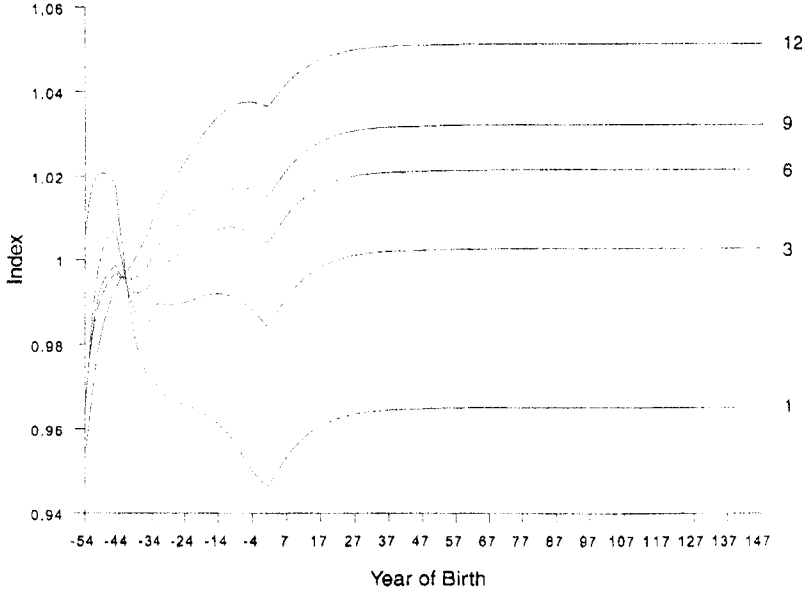

Table 1. Proportional Consumption-Tax Reform

	1996	1997	1998	1999	2000	2005	2010	2025	2055	2145
Aggregates										
National Income[1]	1.000	1.068	1.068	1.070	1.073	1.085	1.094	1.103	1.106	1.106
Capital Stock[1]	1.000	1.023	1.045	1.065	1.084	1.162	1.216	1.289	1.311	1.311
Labor Supply[1]	1.000	1.084	1.076	1.072	1.069	1.061	1.056	1.048	1.045	1.045
Net Saving Rate	0.051	0.102	0.100	0.098	0.095	0.085	0.077	0.065	0.061	0.061
Wage Rates, Interest Rates, and Asset Values										
Before-Tax Wage[1]	1.000	0.986	0.993	0.998	1.003	1.023	1.036	1.053	1.058	1.058
After-Tax Wage[2]	0.775	0.818	0.817	0.821	0.826	0.850	0.866	0.888	0.894	0.894
Before-Tax Rate of Return	0.097	0.102	0.100	0.098	0.096	0.091	0.088	0.083	0.082	0.082
After-Tax Rate of Return	0.080	0.102	0.100	0.098	0.096	0.091	0.088	0.083	0.082	0.082
Normalized Tobin's Q	1.000	0.866	0.859	0.858	0.859	0.867	0.873	0.880	0.882	0.882
Dynamic Replacement Income Tax Rate and Payroll Tax Rate										
Tax Rate[3]	n/a	0.140	0.147	0.148	0.147	0.139	0.134	0.127	0.125	0.125
Payroll Tax Rate	0.147	0.138	0.138	0.138	0.138	0.138	0.139	0.140	0.141	0.141

[1] These variables grow without bound along the balanced-path equilibrium and are represented per-effective labor unit. Hence, they remain constant in the baseline steady state. These variables are indexed with a baseline value of 1.00 in 1996.
[2] The After-Tax Wage rate is computed as $(1-\tau)\times$(Before-Tax Wage) where τ is the economy-wide effective average marginal tax rate on wage income.
[3] Statutory rate, post evasion. Effective rate is about 0.01 smaller due to evasion. Cash flow on capital income also equals this value.

Social Security's Privatization

The U.S. Social Security System faces a grave long-term financial crisis, the full dimension of which is not well known. Paying out benefits on an ongoing basis requires an immediate and permanent increase of roughly 50% in the OASDI payroll tax rate. The President and Congress are now embarked in a national debate about how to save Social Security. They are considering options that include cutting benefits, raising the payroll tax, and privatizing all or part of the system by allowing people to contribute to individual accounts. A key set of issues in this debate is how any policy, including maintaining the status quo, will affect the macro economy as well as rich and poor members of current and future generations. The key finding of my joint research with Smetters and Walliser is that privatizing Social Security can produce a significant long-run improvement in the economy's capital stock and level of per capita output, but that the transitions to such an improved economy are quite long and that how one finances the transition will determine the duration of the transition as well as the

economy's final resting place. Consumption-tax financed privatizations of Social Security produce the fastest transitions and are most efficient. In addition to eliminating a distortionary payroll tax, a consumption tax entails an implicit one-time wealth tax (in taxing retirees when they spend their wealth on consumption) which is nondistortionary. In forcing the elderly to contribute toward paying off the liabilities of the old system, the consumption tax also redistributes from older spenders to younger savers. This promotes national saving. In contrast to consumption-tax finance, income tax finance entails highly distortionary tax rates in the short run that dramatically slows down the transition.

Another of our findings is that using deficit finance to limit the imposition of higher marginal income tax rates as one is paying off Social Security's accrued liabilities makes the economy's short run look better and its long run look worse. Indeed, sufficient reliance on deficit finance will leave the economy in a worse position in the long run than the one at which it started. We've also learned that Social Security's privatization is remarkably progressive with respect to its impact on the lifetime rich and poor alive in the long run. The lifetime poor enjoy a larger improvement in their well being than do the lifetime rich because a larger fraction of the former group's labor earnings would otherwise be subject to payroll taxation.

Figure 2, taken from Kotlikoff, Smetters, and Walliser [1998c], illustrates the A-K OLG Model's analysis of the effects of social security's privatization. The figure considers privatizing the U.S. system and financing the transition with a consumption tax. The figure's top panel shows long-run increases of 40% and 13% in the economy's capital stock and output, respectively. These gains are significant, but they aren't enormous. Moreover, the half-life of this policy is about 30 years.

The Figure's bottom panel shows that the long-run gains are not free. They come at the price of lower utility to initial older and middle-aged generations. Interestingly, all those alive in the long-run, including the richest (group 12) and poorest (group 1) agents, are better off. Since the system being privatized features a highly progressive benefit schedule, but also a highly regressive tax schedule, this finding shows that eliminating the regressive social security payroll tax is more important to the poor than losing the progressive benefit schedule.

Figure 2: Privatizing social Security with Consumption Tax Transition Finance: Macro Effect

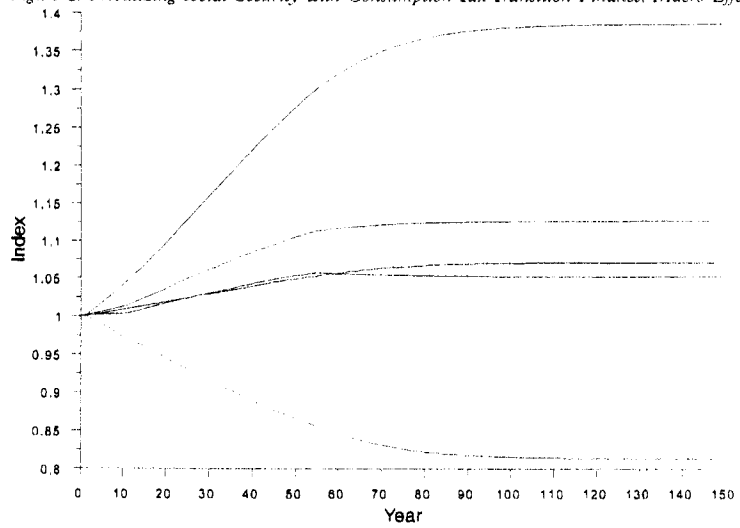

Figure 2: Privatizing Social Security with Consumption-Tax Transition Finance: Remaining Lifetime Utility

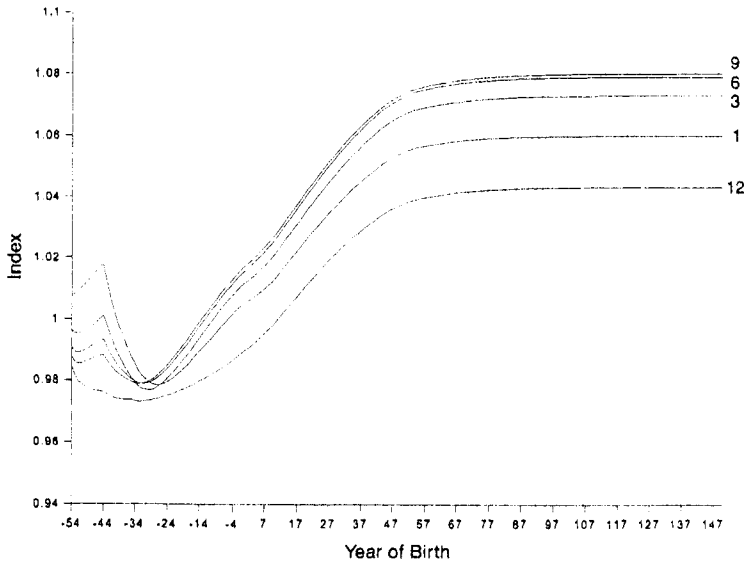

5. Future Improvements to the A-K OLG Model

Demographics

Although Alan and I added demographics to a simplified version of the A-K OLG Model, we are now collaborating with Smetters and Walliser to add demographics to the main A-K OLG Model. The code redesign will be heavily influenced by my previous work on demographics with Alan. Formally, parents are assumed to have utility for (have a utility function defined over) their children's consumption prior to age 21. Each agent gives birth to a specified number of children each year from age 21 through age 45. The numbers of these children are all fractions. When the children reach adulthood (age 21), they, in turn, start giving birth to children. The number of children being born to successive cohorts can be specified exogenously. By specifying a sharp rise and then a sharp decline in this number, one can produce a baby boom followed by a baby bust.

In the present version of the A-K OLG Model, which incorporates intragenerational heterogeneity, adding demographics entails specifying separate birth rates for each earnings group within each cohort. In so doing, we'll have a much richer demographic structure — one that permits different birth rates for the rich and the poor and also permits simulating different changes in birth rates over time for the rich and poor. Thus, for example, we'll be able to simulate a baby bust among high earners and a baby boom among low earners. The present A-K OLG Model also permits one to specify longer lifespans for successive cohorts, so we'll also be able to consider how this factor will impact on the Social Security System.

Immigration

In the course of adding demographics we'll also be able to incorporate immigration into the model. Immigration will be treated as exogenous. The model will be set up so that users can input the number and age structure of immigrations in current and future years. In setting up the model in this manner, we'll be able to consider the potential for immigration policy to improve or exacerbate Social Security's long-term financial condition.

Permitting the Model to Start from an Arbitrary Set of Economic Conditions

In addition to incorporating demographics, a key improvement to the model now underway is permitting it to begin a simulation from an arbitrary set of economic circumstances. At the moment the model solves for an initial steady state and starts policy-reform simulations from that steady state. This is unsatisfactory because there is no reason to think the U.S. economy is currently in a steady state, either with respect to its demographics or its economic fundamentals.

The reason the A-K OLG Model was originally designed to start in a steady state is that the assumptions that the economy is initially in a steady state permit one to calculate the distribution of wealth across and within cohorts. The distribution of wealth provides the critical *state variables* needed to start any simulation. Abandoning the assumption that the economy is initially in a steady state means that one needs to collect data on the intra- and intercohort wealth distribution to provide the starting values of these state variables. We plan to collect this wealth distribution data from the Survey of Consumer Finances (SCF) of the early 1990s. The challenge here will be to find correlates with lifetime labor income (such as education and current labor earnings) since the SCF doesn't permit us to classify wealth by lifetime labor income.

Incorporating Dynamic Programming and Liquidity Constraints

A third improvement on which we're working is to introduce a dynamic programming method of solving for agents' annual levels of consumption, labor supply, and bequests. As we've added more detail to the model to emulate more closely U.S. fiscal institutions, we've stretched our ability to use first-order conditions and shadow prices to solve agents utility maximization problems. The reason is that U.S. fiscal institutions produce highly kinked budget constraints.

Using dynamic programming to solve for the model's micro behavior will also permit us to build on the fundamental work of Hubbard and Judd [1987] in incorporating liquidity constraints in agents' decision making. Liquidity constraints refer to limitations on agents' ability to borrow. Liquidity constraints play an important role in considering certain Social Security privatization proposals. Take proposals that pay off the unfunded liability of the current system by levying a consumption tax. Such proposals would seem to benefit workers because they entail the elimination of the payroll tax and the substitution of a consumption tax that retirees would pay as well as workers. But such proposals also entail compulsory contributions by workers to private saving accounts. Those workers who are liquidity constrained will find that even though they've effectively received a tax break, the requirement of contributing to a new saving account means, in light of their inability to borrow, being forced to lower their current consumption. Whether such workers would actually end up, on balance, worse off is one of the research questions to be addressed in this project. It's possible that, in terms of their remaining lifetime utility levels, their reduced consumption in the short run would be more than offset by higher consumption in the long run.

Adding Money

Adding money to the A-K OLG Model would represent a very important addition. Money creation and the seignorage associated with it is a potentially very important form of financing government's activities and thus represents an important branch of fiscal policy. As Hamann's [1992] work on a simple version of our model showed, one can use the model to study how alternative combinations of monetary and fiscal policy will affect the economy's real variables, as well as its initial price level and time path of inflation.

Adding Housing and Other Goods

A longer term goal is to add housing and other final goods as well as intermediate goods to our model along the lines of Fullerton and Rogers [1993] work. Such an addition would permit us to consider differential commodity taxation as well as differential rates of taxation by industry of the use of capital and labor. Some of the goods that would be added would actually be bads, specifically pollution from the use of fossil fuels. With such an extended model, we could address a number of issues involving environmental policy along the lines of Bohringer, et al. [2000].

Adding Additional Countries

Currently, the A-K OLG Model can be run either as a completely closed or a completely open economy. An obvious extension to the model that would produce results intermediate to those arising under either extreme assumption is to add one or more countries to the model. Indeed, adding a large number of countries would permit the consideration of trade policy as well as monetary and fiscal policy integration.

Adding Idiosyncratic Uncertainty

A final area for future work is adding idiosyncratic uncertainty to the model along the lines of the impressive work of Hubbard, Skinner, and Zeldes [1994a, 1994b, 1995], Imrohoroglu, et al. [1995, 1998a, 1998b], and Huang, et al. [1997]. Three types of uncertainty immediately come to mind: wage rate uncertainty, lifespan uncertainty, and workspan uncertainty. The introduction of dynamic programming to the model is, of course, critically important for addressing these issues. By making uncertainty idiosyncratic, one avoids aggregate uncertainty which would dramatically increase the dimension of the computational problem.

6. Conclusion

The A-K OLG Model has come a long way from the simple prototype Alan Auerbach and I developed in 1980 in my apartment on a machine that would now appear in a computer museum. The model now features intragenerational heterogeneity, does its work in minutes rather than hours, and will shortly feature demographics, begin its simulations from non steady-state positions, and use dynamic programming in solving for micro behavior. Other improvements, like adding money, multiple goods, and idiosyncratic uncertainty, are planned for the near term.

Each improvement increases the model's capacity to study economic policy. But it also makes the model more complex and more of a "black box," particularly to those who are unfamiliar with the various ways of checking the model's specification and predictions. There is no way around this problem except for those using the model to raise comfort levels by providing more detail on their specification checks and as much intuition as possible about their results. The alternative — using very simple models that can be manipulated with the calculus — is helpful in building intuition about the sign of various policy effects, but unsatisfactory for obtaining a quantitative assessment of how and when actual policy changes will affect actual economies.

The other major challenge for the A-K OLG Model, besides improving its capacities, is to bring it "in house" to policy-making institutions around the world. As mentioned, the model is being used by the Congressional Budget Office. It has also been used by economists at the World Bank and the OECD. But the model can and should be used much more routinely and much more broadly in the U.S. and abroad. It can and should be used not just by fiscal authorities, but by monetary authorities as well to study how their policies will affect inflation and alleviate or tighten fiscal constraints. The systematic use of the model and incorporation of its findings into policy debates will force policy makers to confront the long-term implications of their decisions, distinguish cosmetic from real reform, and acknowledge that genuine long-term economic gain generally requires genuine short-term economic sacrifice.

References

Altig, David, Alan J. Auerbach, Laurence J. Kotlikoff, Kent Smetters, Jan Walliser, "Simulating U.S. Tax Reform," *NBER Working Paper No. 6246*, 1997.

Altig, David and Charles T. Carlstrom, "Marginal Tax Rates and Income Inequality in a Life-Cycle Model," *Working Paper*, Federal Reserve Bank of Cleveland, June 1996.

Arrau, Patricio and Klaus Schmidt-Hebbel, "Macroeconomic and Intergenerational Welfare Effects of a Transition from Pay-as-You-Go to Fully Funded Pensions," *Working Paper*, The World Bank, 1993.

Auerbach, Alan J., "Tax Reform, Capital Allocation, Efficiency, and Growth," in Henry Aaron and William B. Gale (eds.) *Economic Effects of Fundamental Tax Reform* (Washington, D.C.: The Brookings Institution, 1996).

Auerbach, Alan J., Robert Hagemann, Laurence J. Kotlikoff, and Giuseppe Nicoletti, "The Economics of Aging Populations: The Case of Four OECD Economies," *OECD Staff Papers*, 1989.

Kotlikoff, Laurence J., Kent Smetters, and Jan Walliser, "Social Security: Privatization and Progressivity," *American Economic Review*, May 1998b.

Kotlikoff, Laurence J., Kent Smetters, and Jan Walliser, "The Economic Impact of Transiting to a Privatized Social Security System," in Horst Siebert (ed.) *Redesigning Social Security* (Kiel University Press, 1998c).

Laitner, John, "Transition Time Paths for Overlapping-Generation Models," *Journal of Economic Dynamics and Control*, 1984.

Lau, Morten I, "Assessing Tax Reforms When Human Capital is Endogenous," in G. W. Harrison, S. E. H. Jensen, L. H. Pedersen, and T. F. Rutherford (eds.) *Using Dynamic Computable General Equilibrium Models for Policy Analysis* (Amsterdam: North Holland, 2000).

Lucas, Robert E., *Supply Side Economics* (Oxford University Press, 1990).

Merette, Marcel, "The Effects of Debt Reduction on Intergenerational Equity and Growth," in Miles Corak (ed.) *Government Finances and Generational Equity* (Ottawa: Statistics Canada, 1998).

Miller, Merton and Charles Upton, *Macroeconomics: A Neoclassical Introduction* (Homewood, IL: R.D. Irwin, 1974).

Raffelhüschen, Bernd, "Anreizwirkungen des Systems der sozialen Alterssicherung. Eine dynamische Simulationsanalyse," *mimeo*, 1989.

Raffelhüschen, Bernd, "Funding Social Security Through Pareto-Optimal Conversion Policies," *Journal of Economics/Zeitschrift fur Nationalokonomie*, Suppl. 7, 1993, 105-31.

Rutherford, Thomas F. "Carbon Abatement in Denmark. Technical Change and Intergenerational Burden Sharing," in G.W. Harrison, S. E. H. Jensen, L. H. Pedersen, and T. F. Rutherford (eds.) *Using Dynamic Computable General Equilibrium Models for Policy Analysis* (Amsterdam: North Holland, 2000).

Schmidt-Hebbel, Klaus. "Pension Reform, Informal Markets, and Long-term Income and Welfare," in Klaus Schmidt-Hebbel (ed.) *Pension Systems: From Crisis to Reform* (Washington D.C.: World Bank, 1999).

Schneider, Ondrej, "Dynamic Simulation of Pension Reform," *mimeo*, Charles University, CERGE-EI, Prague, 1997.

Seidman, Laurence S., "Taxes in a Life Cycle Model with Bequests and Inheritances," *American Economic Review*, 73(3), June 1983.

Shoven, John B. and John Whalley, "A General Equilibrium Calculation of the Effects of Differential Taxation of Income from Capital in the U.S.," *Journal of Public Economics*, 1, 1972.

Slemrod, Joel, and Jon Bakija, *Taxing Ourselves: A Citizen's Guide to the Great Debate over Tax Reform* (Cambridge, MA: MIT Press, 1996).

Summers, Lawrence H, "Capital Taxation and Accumulation in a Life Cycle Growth Model," *American Economic Review*, 71(4), Sept. 1981, 533-44.

Appendix:

Description of the A-K OLG Model

This appendix provides an outline of the current A-K OLG Model, its calibration to the 1996 US economy, and its solution method.

Demographic Structure

The model's cohorts are distinguished by their dates of birth and their lifetime labor-productivity endowments. Following Fullerton and Rogers [1993], each cohort includes 12 lifetime-earnings groups.[4] Each of these 12 groups has its own initial endowment of human capital and its own pattern of growth in this endowment over its lifetime. The lifetime-earnings groups also differ with respect to their bequest preferences. All agents live for 55 periods with certainty (corresponding to adult ages 21 through 75), and each j-type generation is $1+n$ times larger than its predecessor. At model age 21, each j-type cohort gives birth to a cohort of the same type. Population growth is exogenous, and each cohort is $(1+n)^{20}$ larger than its parent cohort.

Preferences and Household Budget Constraints

Each j-type agent who begins her economic life at date t chooses perfect-foresight consumption paths (c), leisure paths (l), and intergenerational transfers (b) to maximize a time-separable utility function of the form

$$U_t^j = \frac{1}{1-\frac{1}{\gamma}}\left[\sum_{s=21}^{75}\beta^{s-21}\left(c_{s,t+s-21}^{j\,\frac{\rho-1}{\rho}}+\alpha l_{s,t+s-21}^{j\,\frac{\rho-1}{\rho}}\right)^{\frac{1-\frac{1}{\gamma}}{1-\frac{1}{\rho}}}+\beta^{54}\mu^j b_{s,t+54}^{j\,\frac{1-\frac{1}{\gamma}}{1-\frac{1}{\rho}}}\right] \quad (1)$$

In (1) α is the utility weight on leisure, γ is the intertemporal elasticity of substitution in the leisure/consumption composite, and ρ is the intratemporal elasticity of substitution between consumption and leisure. The parameter μ^j is a j-type specific utility weight placed on bequests left to each child when the agent dies. The term $\beta = 1/(1+\delta)$ where δ is the rate of time preference, assumed to be the same for all agents.

Letting $a_{s,t}^j$ be capital holdings for type j agents, of age s, at time t, maximization of (1) is subject to a sequence of budget constraints given by

$$a_{s+1,t+1}^j = (1+r_t)(a_{s,t}^j + g_{s,t}^j) + w_{s,t}^j(E_{s,t}^j l_{s,t}^j) - c_{s,t}^j - \sum_{k\in\tilde{T}} T^k\left(B_{s,t}^{j,k}\right) - Nb_{s,t}^j \quad (2)$$

$$l_{s,t}^j \le E_{s,t}^j$$

where r_t is the pretax return to savings, $g_{s,t}^j$ are gifts received from parents, $E_{s,t}^j$ is the time endowment, $b_{s,t}^j$ denotes bequests made to each of the $N = (1+n)^{20}$ children, and the functions $T^k(\cdot)$ with tax base arguments $B_{s,t}^{j,k}$ determine net tax payments from income sources $k \in \tilde{T} = \{C, K, W, Y, P\}$. $T^C(\cdot)$, $T^K(\cdot)$, $T^W(\cdot)$, $T^Y(\cdot)$ and $T^P(\cdot)$ are consumption taxes, capital income taxes, wage taxes, income taxes and social security payroll taxes, respectively. Social security benefits are represented in equation (2) as negative taxes with the base switching at the point of retirement from the contemporaneous payroll base to average indexed yearly earnings in the pre-retirement years. All taxes are collected at the household level and the tax system includes both a personal income tax and a business profits tax. The bases for the wage and payroll taxes are smaller than total labor income due to the base reductions discussed below.

An individual's earnings ability is an exogenous function of her age, her type, and the level of labor-augmenting technical progress, which grows at a constant rate λ. We concentrate all skill differences by age and type in an efficiency parameter ε_s^j. Thus, the wage rate for an agent of type j and age s is $w_{s,t}^j = \varepsilon_s^j w_t$, where w_t is the growth-adjusted real wage at time t. ε_s^j increases with age to reflect not only the accumulation of human capital, but also technical progress. To permit balanced growth for our

specifications of preferences given the restriction on leisure shown in equation (2), we assume that technical progress also causes the time endowment of each successive generation to grow at rate λ.[5] Thus, if $E_{s,t}^j$ is the endowment of type j at age s and time t, then $E_{s,t}^j = (1+\lambda)\,E_{s,t-1}^j$, for all s, t, and j. Notice that the endowment $E_{s,t}^j$ depends only on an agent's year of birth. Because E grows at rate λ from one cohort to the next, there will be no underlying trend in w_t. The growth-adjusted earnings ability profiles take the form

$$\varepsilon_s^j = e^{a_0^j + a_1^j s + a_2^j s^2 + a_3^j s^3}. \quad (3)$$

Values of the a coefficients for j-type groups 1 through 12—in ascending order of lifetime income — are based on regressions fitted to the University of Michigan's Panel Study of Income Dynamics and are taken from Altig, Auerbach, Kotlikoff, Smetters and Walliser [1997]. Groups 1 and 12 comprise the bottom and top 2% of lifetime wage income earners, and groups 2 and 11 the remaining 8% of the top and bottom deciles. All other groups constitute 10% of the population. For example, group 3 is the second decile of lifetime-wage income, group four the third decile, and so on up to group 10. The estimated earnings-ability profiles, scaled to include the effects of technical progress. Given our benchmark parameterization, peak hourly wages valued in 1996 dollars are \$4.00, \$14.70, and \$79.50 for individuals in classes 1, 6, and 12, respectively. More generally, steady-state annual labor incomes derived from the model's assumptions and the endogenous labor supply choices range from \$9,000 to \$130,000. These calculations do yet include labor compensation in the form of fringe benefits (discussed below).

Transfers are received by children, with interest, at the beginning of the period after they are made by their parents. We restrict all parental transfers to bequests, so that $b_{s,t}^j = 0$, for $s \neq 75$, and $g_{s,t}^j = 0$, for $s \neq 56$. In the steady state, therefore, $g^j = b^j$, for all j (where we have dropped the age subscripts for convenience). The parameters μ^j are derived endogenously for the initial steady state such that the ratio of the bequest to economy-wide mean income corresponds to the ratio originally estimated by Menchik and David [1982] and updated by Fullerton and Rogers [1993]. Bequests range from \$4,800 to \$450,000 for the lowest and highest lifetime earnings classes, respectively.

Choices for the remaining technology, preference, and demographic parameters are summarized in Table 1. The benchmark values for δ, γ, ρ, and n are those in Auerbach and Kotlikoff [1987]. The parameter α is

chosen so that agents devote, on average, about 40% of their available time endowment (of 16 hours per day) to labor during their prime working years (real-life ages of roughly 21-55).

The Non-Social Security Government Budget Constraint

At each time t, the government collects tax revenues and issues debt (D_{t+1}) which it uses to finance government purchases of goods and services (G_t) and interest payments on the inherited stock of debt (D_t). Letting ϕ^j be the fraction of j-type agents in each generation, the non-social security part of the government's budget constraint evolves according to

$$D_{t+1} + (1+n)^t \sum_{j=1}^{12} \phi^j \sum_{s=21}^{75} (1+n)^{-(s-21)} \sum_{k \in \{T-P\}^j} T^k\left(B_{s,t}^{j,k}\right) = G_t + (1+R_t) D_t \quad (4)$$

The exclusion of social security taxes in equation (4) reflects the fact that social security currently uses self-financing earmarked taxes.

Government expenditures are assumed to be unproductive and generate no utility to households.[6] The values of G_t and D_t are held fixed per effective worker throughout the transition path. Any reduction in government outlays resulting from a change in the government's real interest payments is passed on to households in the form of a lower tax rate. The level of government debt, D_t, was chosen such that the associated real interest payments equal about 3.5% of national income in the initial steady state. The statutory tax schedules (described below) generate a level of revenue above debt service such that the benchmark steady-state ratio of government purchases, G_t, to national income equals 0.239. These values correspond very closely to the corresponding 1996 values for the combined local, state, and federal government in the United States. See Table 2.

Non-Social Security Taxes

The benchmark tax system in our initial steady state is designed to approximate the salient aspects of the 1996 U.S. (federal, state, and local) tax and transfer system. It features a hybrid tax system (incorporating wage-income, capital-income, and consumption tax elements) and payroll taxation for the Social Security and Medicare programs. To adjust for tax evasion, we reduce income taxes by 2.6 percentage points. This adjustment

is consistent with the degree of tax evasion reported in Slemrod and Bakija [1996]. In the various alternative tax structure experiments we assume that evasion reduces the post-reform tax base (income net of deductions and exemptions) by the same percentage as before the reform. Thus, the level of tax evasion falls when the tax base shrinks.

We approximate the hybrid current U.S. tax system by specifying a progressive wage-income tax, a flat capital-income tax, a flat state income tax, and a flat consumption tax.

Wage Income Taxation

The wage-income tax structure has four elements: 1) a progressive marginal rate structure derived from a quadratic approximation to the 1996 federal statutory tax rates for individuals, 2) a standard deduction of $4000 and exemptions of $5660 (which assumes 1.2 children per agent, consistent with the model's population growth assumption), 3) Itemized deductions—applied only when they exceed the amount of the standard deduction—that are a positive linear function of income estimated from data reported in the *Statistics of Income*,[7] and 4) Earnings-ability profiles that are scaled up to incorporate pension and non-pension components of labor compensation.[8]

The model's initial economy-wide average marginal tax rate on wage income is about 21%, about the figure obtained from the NBER's TAXSIM model reported in Auerbach [1996]. The average wage-income tax rate equals 12.1%. For all individuals in the highest lifetime income class (group 12), the average effective marginal tax rate on labor income is 28.6%. The highest realized effective marginal tax rate is 34%. For lifetime income class 6—whose members have peak labor earnings of about $35,000—the average tax rate and average marginal tax rate are 10.6 and 20.0%, respectively. For the poorest class (group 1), the corresponding rates are zero and 5.5%.[9]

Capital Income Taxation

Following Auerbach [1996], we assume that income from residential capital and non-residential capital are taxed at flat rates of 6% and 26%, respectively. Given the roughly equal amounts of these two forms of capital, the effective federal marginal tax rate on total capital income is 16%. However, this rate

applies only to new capital. Existing capital faces a higher tax rate which, given depreciation schedules, is estimated to be 20%. We model this gap by assuming that all capital income faces a 20% tax, but that 20% of new capital may be expensed, thereby generating a 16% effective rate on new capital.

State Income Taxation

In addition to the federal taxation, both capital and wage income are subject to a proportional state income tax of 3.7%. This value corresponds to the amount of revenue generated by state income taxes in 1996 divided by national income.

Consumption Taxation

Consumption taxes in the initial steady state reflect two elements of the existing tax structure. First we impose an 8.8% tax on consumption expenditures consistent with values reported in the National Income and Product Accounts on indirect business and excise revenues. However, because contributions to both defined benefit and defined contribution pension plans receive consumption tax treatment, we levy an additional 2.5% tax on household consumption goods expenditures to account for the indirect taxation of labor compensation in the form of pension benefits (Auerbach [1996]). This 2.5% tax replaces the wage tax that otherwise would apply to labor compensation in the form of fringe benefits.

Social Security, Medicare and Disability

The model has a social insurance system that incorporates social security Old-Age and Survivors Insurance (OASI), Social Security Disability Insurance (DI), and public health insurance taking the form of Medicare (HI).

OASI benefits are calculated according to the progressive statutory bend-point formula. U.S. Social Security benefits are based on a measure of average indexed monthly earnings (AIME) over a 35-year work history. The AIME is converted into a primary insurance amount (PIA) in accordance with a progressive formula. In particular, the 1996 benefit

formula has two bend points. The PIA is calculated as 90% of the first $437 of AIME, 32% of the next $2,198 of AIME, and 15% of AIME above $2,198. We approximate the benefit formula with a sixth-order polynomial which is applied to the dollar-scaled AIME generated by the model. This polynomial approximation is very accurate with a $R^2 = 0.99$ (Figure 1). We achieve replacement values between 25 and 75% for the lifetime richest and lifetime poorest, respectively. Since approximately 50% of Social Security benefits are paid to survivors and spouses, we multiply benefits by a factor of two.

An earmarked tax applied to wage income up to a limit of $62,700—the earnings ceiling in 1996—is used to pay for OASI benefits. Define $\omega^j_{s,t} \equiv w_{s,t}(E^j_{s,t}l^j_{s,t})$ as the wage income earned by the j-type agent who is age s in year t. Also define $\bar{\omega}^j_{65,t}$ as the average indexed annual earnings for the j-type agent age 65 at time t. Labor income earned before turning age 65 is adjusted upward by the growth rate of the economy in calculating $\bar{\omega}^j_{65,t}$. Payroll taxes at time t—with retirement benefits modeled as negative taxes—equals

$$T^P(B^{j,k}_{s,t}) = \begin{cases} \tau \cdot \omega^j_{s,t} & ; \; s \le 64 \,, \omega^j_{s,t} \le \$62{,}700 \\ \tau \cdot \$62{,}700 & ; \; s \le 64 \,, \omega^j_{s,t} > \$62{,}700 \\ -2 \cdot R(\bar{\omega}^j_{s,t}) \cdot \bar{\omega}^j_{s,t} & ; \; s > 64 \end{cases} \tag{5}$$

where R(·) is the statutory replacement rate function shown in Figure 1.

Budget balance for a self-financing pay-as-you-go social security system with earmarked taxes at time t requires:

$$\sum_{j=1}^{J} \phi^j \sum_{s=21}^{75} (1+n)^{-(s-21)} \, T^P(B^{j,P}_{s,t}) = 0 \tag{6}$$

The value of τ is solved for endogenously as a function of benefit rules via equation (6). The value of $\hat{o}$ is 9.9% in the initial steady state, which is close to its actual value in 1996.[10]

The net marginal tax rate is a component of the consumer's first-order conditions. Let $PVT(\omega^j_{s,t})$ and $PVB(\omega^j_{s,t})$ be the present value of payroll

taxes and benefits, respectively, for the j-type agent age s at time t. The net marginal tax rate for those below the earnings ceiling in each case considered herein is:

$$\theta(\omega^j_{s,t}) = \begin{cases} \tau \cdot \left[1 - PVB'(\omega^j_{s,t})/PVT'(\omega^j_{s,t})\right] & ; \textit{full perception linkage} \\ \tau & ; \textit{no perception linkage} \end{cases} \quad (7)$$

where $PVB'(\cdot) = \partial\ PVB(\cdot)/\partial\omega$ and $PVT'(\cdot) = \partial\ PVT(\cdot)/\partial\omega$. The net marginal tax rates under the perception linkage are shown in Figure 2 by income class and age. These tax rates are typically relatively higher for both richer and younger agents. The higher rates for richer agents reflect the progressive manner in which social security benefits are calculated. The higher rates for younger agents reflect the compound interest effect of being required to save in a social security system whose internal rate of return is less than after-tax rate of return to capital (reported below). Notice that the net tax rates are generally quite large and positive even for the lifetime poor because the after-tax rate of return to capital is higher than the internal rates of return faced by these agents. Rich agents whose labor income exceeds the payroll tax (e.g., class 12 in select years) face a zero marginal tax rate.

The HI and DI programs are modeled very simply. The HI and DI levels of lump-sum transfers are picked to generate payroll tax rates of 2.9% and 1.9%, respectively, corresponding to their 1996 statutory rates. Like the OASI tax, DI contributions apply only to wages below $62,700. The HI tax, in contrast, is not subject to an earnings ceiling. Lump-sum HI and DI benefits are provided on an equal basis to agents above and below age 65, respectively.

Aggregation and Technology

Aggregate capital (K) and labor (L) are obtained from individual asset and labor supplies as

$$K_t = (1+n)^t \sum_{j=1}^{12} \phi^j \sum_{s=21}^{75} (1+n)^{-(s-21)} a^j_{s,t}\ D_t, \quad (8)$$

(where, recall, D_t is government debt at time t) and

$$L_t = (1+n)^t (1+\lambda)^{-t} \sum_{j=1}^{12} \phi^j \sum_{s=21}^{75} (1+n)^{-(s-21)} e^j_s (E^j_{st} - l^j_{st}). \quad (9)$$

Output (net of depreciation) is produced by identical competitive firms using a neoclassical, constant-returns-to-scale production technology. The aggregate production technology is the standard Cobb-Douglas form

$$Y_t = AK_t^{\theta} L_t^{1-\theta}, \quad (10)$$

where Y_t is aggregate output (national income) and θ is capital's share in production. Denote the capital-labor ratio as κ. The time-t competitive post-tax capital rate of return equals

$$r_t = \left[\ \theta A \kappa_t^{\theta-1} (1-\tau^K_t) + q_{t+1} - q_t\right] / q_t. \quad (11)$$

where $q_t = \left(1 - z_t \tau^K_t\right)$ is Tobin's q at time t and z is the level of capital investment expensing.

Given our parameter choices, the model generates a pre-tax interest rate of 9.3%, a net national saving rate of 5.3%, and a capital/national-income ratio of 2.6. Consumption accounts for 73.4% of national income, net investment for 5.2%, and government purchases of goods and services for 21.4%. These figures are close to their respective 1996 NIPA values. The post-tax interest rate equals 0.08 and is calculated following Auerbach [1996].

Solving the Model

The model solves for the full rational-expectations dynamic (Nash) equilibrium with a Gauss-Seidel algorithm. The calculation starts with a guess for certain key variables and then iterates on those variables until a convergence criterion is met. The identifying restrictions of the model are used to compute the remaining economic variables as well as the updates for the iterations. The solution involves several steps and inner loops that solve for household-level variables before moving to an outer loop which solves for the time-paths of aggregate variables and factor prices. Since the decision to opt out by any agent will be affected by the exact time path of factor prices—which, in turn, is affected by the opting out decisions of other agents—the opting out choice is determined endogenously for each agent. The solution algorithm iterates until each agent, given the prevailing path of factor prices, prefers his/her

intertemporal allocation of consumption and leisure and his/her decision whether to opt out.

The household optimization problem is subject to the constraint that leisure not exceed the endowment of time (equation (2)). For those households who would violate the constraint, the model calculates shadow wage rates at which they exactly consume their full-time endowment.

The household's budget constraint is kinked due to the tax deductions applied against wage income. A household with wage income below the deduction level faces marginal and average tax rates equal to zero. A household with wage income above the deduction level faces positive marginal and average tax rates. Due to the discontinuity of the marginal tax rates, it may be optimal for some households to locate exactly at the kink. Our algorithm deals with this problem as follows. We identify households that choose to locate at the kink by evaluating their leisure choice and corresponding wage income above and below the kink. We then calculate a shadow marginal tax rate from the first-order conditions that puts those households exactly at the kink. This procedure generates optimal forward-looking leisure and consumption choices for all periods of life.

The payroll tax ceiling introduces additional complexity by creating a non-convexity in the budget constraint. For those above the payroll tax ceiling, the marginal tax rate on labor falls to zero. We evaluate the utility on both sides of the non-convex section and put households on the side that generates highest utility.

The sequence of calculations is as follows. An initial guess is made for the time-paths of these aggregate variables as well as for the shadow wage rates, shadow tax rates, endogenous tax rates, the separate OASI / DI / HI payroll tax rates, and the Social Security and Medicare wealth levels. The corresponding factor prices are calculated along with the forward-looking consumption, asset and leisure choices for all income classes in each current and future cohort. Shadow wages and shadow taxes are calculated to ensure that the time endowment and the tax constraints discussed above are satisfied. Households' labor supply and assets are then aggregated by both age and lifetime income class at each period in time. This aggregation generates a new guess for the time-paths of the capital stock and labor supply. The tax rate which is endogenous for the particular simulation, is updated to meet the revenue-neutrality requirement. The payroll tax is also updated to preserve the pay-as-you-go financing of OASI and HI benefits.[11] The tax rate for DI benefits is also updated. The algorithm iterates until the capital stock and labor supply time-paths converge.

Notes

* Boston University and The National Bureau of Economic Research. I thank Alan Auerbach for very helpful comments.

[1] Which appeared as a mimeo in 1977.

[2] Which appeared as a mimeo in 1981.

[3] See Gokhale, Kotlikoff, and Sabelhaus [1996] for an empirical study of propensities to consume of U.S. cohorts.

[4] Our model has several strengths relative to Fullerton and Rogers [1993] and, at least, one weakness. The strengths include a rational-expectations solution, a social security system, a tax system with progressive marginal tax rates, an array of tax base reductions, government debt, bequests, and other features. The model herein, however, lacks the multi-sectoral detail on the production side present in the Fullerton-Rogers model. The omission of this production detail probably has little relevance for our purposes since privatization does not change the inter-sectoral distortions.

[5] See Auerbach, et al. [1989] for a more complete discussion of this strategy for dealing with balanced growth.

[6] Since G remains fixed in all of our experiments, incorporating G into the utility function is unimportant.

[7] The data used in this estimation was taken from all taxable returns in tax year 1993. The function was obtained by regressing deductions exclusive of mortgage interest expense on the midpoints of reported income ranges. (The deduction of interest expense on home mortgages was included in our calculation of the capital-income tax rate, as we will subsequently describe.) The regression yielded a coefficient of 0.0755 with an R^2 equal to 0.99.

[8] Benefits as a function of adjusted gross income were kindly provided by Jane Gravelle of the Congressional Research Service and Judy Xanthopoulos of the Joint Committee on Taxation, respectively. Based on this information we regressed total benefits on AGI. The regression yielded a coefficient of 0.11295 with an R^2 equal to 0.99. In defining the wage-tax base, we therefore exempt roughly 11% of labor compensation from the base calculations.

[9] The average marginal rate for people with the lowest income exceeds zero due to positive shadow tax rates in peak earnings years.

[10] The employer-employee combined payroll tax equaled 10.52 percentage points. Approximately, 1 percentage point of this tax represents a net increase to the social security trust fund – an institution we don't model. Hence, the relevant payroll tax for our calibration is 9.52%.

[11] Note that the Social Security replacement rate and absolute level of Medicare benefits are exogenous.