University of Michigan Retirement Research Center

Modeling the Macroeconomic Implications of Social Security Reform

John Laitner

M R R C Issues in Brief

Executive Summary

Two models dominate the economics literature on why people save (and hence on why they accumulate wealth). In one, the *life–cycle* model, people save when they are young in order to have funds to support their retirement. In the other, the *dynastic* model, households save to build estates to pass to their descendants. From the standpoint of public policy analysis, the two models can have quite different predictions. It is also the case that research shows that either model by itself has shortcomings in describing actual savings behavior. In this *Issues in Brief*, I summarize how one might combine the two basic models, what advantages the combined framework might have in matching empirical evidence, and what implications for public policy — specifically Social Security reform — the combined model might have.

Background

The life–cycle model is manifestly attractive. It emphasizes a basic pattern of life: a household's earnings tend to rise during youth and middle age, but they often disappear in old age, with retirement. A household with low earnings might be comfortable retiring on Social Security benefits alone. To maintain reasonably level annual consumption spending, most households, however, must save in youth and middle age to build private wealth, which they can spend in their retirement.

A closer look raises some doubts about the life-cycle model as the sole explanation for savings behavior. Although the model predicts that households will seek to spend their net worth down to zero before death, economists have never been sure empirical evidence supports such complete depletion. Similarly, there has

been longstanding suspicion that the life–cycle model cannot explain the vast total amount of wealth that we see in the U.S. economy (for example, in the Federal Reserve Bank's *Flow of Funds* data). Furthermore, the distribution of wealth holdings in the U.S. is very unequal. For instance, the Federal Reserve Bank's *Survey of Consumer Finances 1995* shows the wealthiest 1% of U.S. households own 30% or more of U.S. private wealth, and the top 5% own over half. This is much more unequal than the distribution of earnings, where the share of the top 1% is about 11%. Simulations of life–cycle models generally fail to generate nearly as much concentration of wealth as these data indicate.

According to the dynastic model, households may save to build estates to pass to their descendants. In the simplest version of that framework, the so–called representative agent model, all households are alike and all leave bequests. Due to the framework's analytic tractability, economists employ it widely. Another advantage of it is that it generally has little difficulty explaining the total amount of private net worth that we see. Nevertheless, in terms of matching survey data, the representative agent model is unsatisfactory. For one thing, only about half of U.S. families report that they receive an inheritance at some point in their lives, and most reported inheritances are quite small. The long time span of the *Panel Study of Income Dynamics* (PSID) data set makes it convenient for studying intergenerational linkages. Yet, recent PSID–based work on the correlation of the consumption of parents and their grown children, on *inter vivos* transfer behavior, and on inheritances, fails to support the simple dynastic model, or provides at most mixed support.

The Model

I propose a model combining life–cycle saving and estate building. Key elements of the combined system are that every household does life–cycle saving; the earnings levels of different households are different; all households care about both their own lifetime well-being and that of their descendants; and, liquidity constraints exclude negative bequests and negative net worth at any age. The model has a simple aggregate production function; labor supply is inelastic; earning ability differences among households are exogenous; and, the economy is closed.

In terms of matching empirical evidence, the hybrid model offers potential advantages over either basic model taken separately. Wealth accumulation will tend to be less equal than in the pure life–cycle model. As in the life cycle model, every household saves during the first part of its life in anticipation of retirement. However, estates are a different matter. Since earning abilities within family lines regress over generations toward the mean, low earners expect that their descendants will have favorable earnings relative to their own. Therefore, they will not attempt to build an estate. High earners face the opposite circumstance: parents with high earnings must expect that their descendants will likely do worse; hence, they have incentives to build estates to share their good luck. In the end, higher earners save for two reasons—to support their retirement, and to leave an estate. Lower earners save for only one reason—to support their retirement. It is also true that the hybrid does not predict that all parents will leave bequests. On the contrary, only parents with high earnings and/or large inheritances will want or be able to do so.

My work calibrates the combined model. I set the distribution of earning abilities from the *Survey of Consumer Finances 1995*. Several other parameters are quite standard. Two are less so —

namely, a parameter determining the degree of household risk aversion, and a parameter determining the importance which parents place on their grown children's lifetime well–being relative to their own. I jointly calibrate the last two parameters so that the model's equilibrium has as much private wealth (relative to gross domestic product) as the U.S. economy, and that the model generates a distribution of bequests yielding realistic federal estate tax revenues. In the best calibrations, households are quite tolerant of risk, and parents weight their grown children's well-being about 80% as high as their own.

Summary of Findings

A simulation based on the best calibration matches several aspects of the U.S. economy surprisingly well. In the simulation, about half of all households in each birth cohort choose to leave bequests. And the simulated distribution of wealth is very unequal. For example, the richest 1% of households in every cross section own about 25% of total private net worth.

Turning to analysis of public policy, the pure life–cycle model suggests that reform instituting funded private accounts for Social Security in place of part of the current unfunded Social Security system could substantially increase the economy's wealth accumulation in the long run. The idea is that taxes reduce household resources for saving, and benefits reduce households' need to save for their retirement. The dynastic model has the opposite implication: since the present value of any system's taxes and benefits exactly balance for each endless dynasty, neither a funded nor an unfunded system affects the representative agent's dynastic well-being or consumption. Private intergenerational transfers counterbalance public programs, and the economy's privately supported capital stock remains always the same.

My work shows the hybrid model can exhibit either reaction to policy, or indeed any reaction between the two extremes. The wide range of potential outcomes makes precise calibration of the model's parameters very important. The best calibrations at this point yield a combined model with policy implications resembling the representative agent dynastic formulation rather than the pure life–cycle case. This is true despite the fact that in the simulation life–cycle saving alone finances about 70% of total private net worth.

Conclusion

Further research is needed. For example, future research will incorporate lifetime earnings uncertainty for households. Perhaps lifetime precautionary saving will further reduce the importance of estate building. Even if present results hold, other analysis suggests that reform might play a role in controlling wealth inequality. Reductions in inequality might be a desirable—though otherwise unexpected—concomitant of funding Social Security or reducing the national debt.

In conclusion, a relatively straightforward combination of existing economic models seems to offer the potential of explaining a number of empirical features of the U.S. economy that have heretofore been somewhat opaque. Precise calibrations of the parameters of the combined model are important since the potential range of policy implications is broad. At this point, analysis suggests that outcomes from Social Security reform such as greater equality, flexibility, and sustainability may be at least as important as potential long–run increases in national saving.

John Laitner is the Director of the Michigan Retirement
Research Center and a Professor of Economics at the University
of Michigan; Ann Arbor, Michigan 48106 (jlaitner@umich.edu).