

What determines savings?

A review essay

David N. Weil*

Brown University, Providence, RI 02912, USA

1. Introduction

What Determines Savings? (MIT Press, Cambridge, 1989) collects 19 essays, most previously published, by Laurence Kotlikoff and an assortment of co-authors.¹ This is a diverse set of work, ranging from extensions of theory to empirical tests to policy simulations. The book's unifying perspective is that the best way to understand saving is to model people as optimizing over a very long horizon, that is, over at least their own lifetimes. Whether the horizon is longer than the lifetime, that is, whether the infinite horizon model is superior to the life cycle model, is a question addressed by several of the essays. Several other essays add different twists to the 'pure' life cycle model – such as uncertainty about lifespan or about health expenditures, or within-family annuities – that cause pure life cycle optimizers to engage in behavior similar to infinite horizon optimizers. And several essays start with the pure life cycle model as a base for simulating the effects of government policy or demographic change.

Overall, the essays in this book are impressive in their thoroughness and originality. Kotlikoff has demonstrated remarkable energy in pursuing research on saving – and he is open-minded enough that several of the essays are in direct contradiction to each other. This can make reading his book either exhilarating or exasperating, depending on one's inclinations.

Although there is a large amount to be learned from this book, I see two problems with the style of research that it represents. The central theoretical

*I am grateful to Chris Carroll, David Cutler, and Louise Sheiner for helpful comments.

¹The co-authors are Alan Auerbach (three times), Michael Boskin (once), John Shoven (once), Avia Spivak (three times), and Lawrence Summers (twice).

tension in the book is between the pure life cycle model and its cousin, the infinite-horizon model. One problem is the failure to recognize that the same model may not apply to all wealth classes. The second problem is the mismatch between theory and the possibility of empirical work. I will address these points in my discussion of the papers presented in the book, and discuss them more fully at the end of this review.

The book begins with a long introductory essay summarizing the state of both theoretical and empirical research in saving. The essay can be used as either an overview for a reader unacquainted with the field, or as a reference for someone currently doing research (to the latter, the book's 24-page bibliography is invaluable).

2. Life cycle theory

Two essays deal with the relation between annuities and saving. Annuities are of great interest to economists studying saving in a life cycle context because they allow individuals to maximize expected utility when facing an unknown date of death. When there is no bequest motive, the increase in expected utility from purchasing an annuity as compared to holding a dwindling stock of assets can be quite large. Since there are good reasons for people to buy annuities, and yet we observe few annuities purchased, annuities play the role of 'the dog that did not bark' in the no-bequest-motive life cycle model [Friedman and Warshawsky (1990)].

One way of getting around this problem is to argue that the elderly do in fact acquire annuities, but that they do not do so in an observable market. 'The Family as an Incomplete Annuities Market' makes the argument that annuities within the family – either within a marriage or between parents and children – can substitute in large part for an external annuities market. The key assumption of the model is that problems of information and enforceability are more easily solved within the family than they are in the market, even though family members are not modelled as being altruistically linked.

The second paper in this series, 'Annuity Insurance, Savings, and Inequality', considers how the introduction of fair market annuities would affect the level of savings for a series of models of family types: with and without altruistic linkages and with and without within-family annuities. Not surprisingly, the biggest effect of fair annuities on saving comes in the model where there are neither altruistic linkages nor within-family annuities. An interesting aspect of this paper is its attempt to link the distribution of wealth in the economy to accumulation of capital by intergenerationally linked families. Given the evidence, described below, that wealth from bequests is an important component of the capital stock, it is natural to ask what other phenomena in the economy may be linked to intergenerational transfers. The

model presented in this paper shares with Abel's (1985) model the property that the distribution of wealth is the product of a kind of mortality lottery. The rich are those whose families have had the 'luck' that all of their recent ancestors died at the beginning of retirement, and so did not spend down their resources.

3. Applying the life cycle model

Several of the essays are concerned with the effects of fiscal policy on saving and on the welfare of different generations. 'Deficit Delusion' argues that the conventionally used definition of the deficit is arbitrary, since the government can relabel streams of receipts as 'borrowing' or 'taxes' and outflows as 'spending', 'interest on loans', or 'repayment of loans'.

Given this critique, the natural question is, how *should* we analyze the intertemporal aspects of fiscal policy? Kotlikoff argues for looking at the impact of policies on the resources of different generations. Thus, for example, policies that transfer huge amounts of wealth between generations, like the 1982 social security reforms, should bulk large in the analysis of fiscal policy and saving, while policies that do not affect the different cohorts' budget constraints should be expected to have little impact. But this approach is really just a restatement of the idea that individuals' lifetimes are the important planning horizon in saving. If the horizon is shorter, because of liquidity constraints or myopia, then the conventionally defined deficit will indeed affect saving. And if the planning horizon is longer than the lives of the current population, then policies that transfer wealth between generations are not important because they will be undone by altruistically linked families.

'Taxation and Savings: A Neoclassical Perspective' analyzes fiscal policy in a world where people see through the intertemporal aspects of fiscal policy, and can freely undo any changes in the timing of their disposable income. In such a world, the response of saving to a cut in taxes will depend in sign and magnitude on how individuals expect the resulting imbalance in the government budget constraint to be corrected – for example, on which generation the burden will fall. Time series regressions that do not control for 'the future time path of accommodating fiscal policy... are of dubious value in describing the potential saving impact of changes in fiscal policy' (p. 191). Since in the real world the government does not announce all of its future fiscal plans, it is hard to know what to make of this sort of analysis.

Several of the essays rely on life cycle simulation models to answer interesting questions about savings. Many of them use variations on the model presented by Auerbach and Kotlikoff in *Dynamic Fiscal Policy* (1987). The Auerbach–Kotlikoff (AK) model finds equilibrium paths for an economy populated by 55 overlapping generations of nonaltruistic agents. Labor sup-

ply, saving, capital stock, wages, and interest rates are endogenous, and all decisions are made by forward-looking agents in the absence of uncertainty. For example, 'Investment versus Savings Incentives: The Size of the Bang for the Buck and the Potential for Self-Financing Business Tax Cuts' uses the AK model to analyze the effects of savings policies (which treat all capital equally) versus investment policies (which treat newly produced capital more favorably than old capital).

In an interesting blend between theory and empirical work, 'An Examination of Empirical Tests of Social Security and Savings' uses the AK model to demonstrate graphically the misspecification of conventional time series tests of the effects of social security. Unfunded social security is introduced into an economy that had been in steady state, and transition paths of all of the endogenous variables are calculated. Time series regressions of consumption on income, wealth, and social security wealth are then run using the data generated by the model. The size and sign of the coefficient on social security wealth turn out to be extremely sensitive to the sample period over which the regression is run.

4. Tests of the life cycle model

The most interesting papers in the book are those that confront the life cycle model with the data. 'The Role of Intergenerational Transfers in Aggregate Capital Accumulation' addresses the ability of the life cycle model to explain aggregate savings. The paper presents a framework in which all wealth can be divided into two types: Life cycle wealth is the difference between the present values of past earnings and consumption of people now alive. Transfer wealth is the present value of net intergenerational transfers received by people now alive. Since transfer wealth is hard to observe directly, the relative importance of the two sources of wealth must be ascertained by comparing total wealth to life cycle wealth.

Calculating life cycle wealth requires knowing the historical paths of earnings and consumption for all individuals currently alive. To construct these, the paper uses cross-sectional estimates of relative consumption and earnings of people of different ages, along with time series on aggregate earnings and consumption. It is evident from examining the earnings and consumption profiles calculated by this method that the life cycle model is in trouble. Earnings and consumption for the members of a given cohort are almost equal through age 50. From age 50 through 65 labor income is slightly higher than consumption, and after age 65 earnings fall below consumption. Even just before retirement, however, individuals' life cycle wealth is small.

Accumulating past differences between earnings and consumption for people currently alive using a variety of interest rates produces a striking

result: life cycle wealth accounts for only a small fraction (at most 20%) of the net worth of the noninstitutional household sector.

The second question asked in this paper is how large the capital stock would be in the absence of intergenerational transfers. As long as the interest rate is larger than the growth rate of the economy – so that the bequest that an individual receives is larger in present value than the bequest that he will leave – the presence of bequests pushes out an individual's budget constraint. Eliminating bequests would induce people to work more or consume less. Thus the decrease in wealth contingent on the elimination of bequests would be smaller than the current stock of transfer wealth.² But this effect does not vitiate the paper's main result on the importance of transfer wealth: for various life cycle utility functions, the paper calculates the reduction in wealth that would come from eliminating bequests as being around 70% of the current stock of transfer wealth.

It should be noted that if people are not life cycle savers in the sense that they solve a lifetime optimization problem without capital market imperfections, the elimination of transfers may have a much smaller impact on wealth. Consider a world in which people receive intergenerational transfers early in life, spend down their assets to zero at some point in the middle of life, and then accumulate assets before they die. If most of the wealth in this economy is held by people who are beyond their mid-life liquidity crises, then eliminating intergenerational transfers will have little effect on aggregate wealth stock, even though the accounting scheme of the paper might label all wealth as transfer wealth. A similar example can be constructed in a world where people are 'buffer-stock' savers described below.

'Intergenerational Transfers and Savings', which appeared alongside Franco Modigliani's defense of the Life Cycle Model in the *Journal of Economic Perspectives*, presses the argument that intergenerational transfers account for most of the stock of wealth. One component of intergenerational transfers is *inter-vivos* gifts, and part of the debate between Kotlikoff and Modigliani is over what expenditures on children should be considered as part of parents' consumption and which should be considered as intergenerational transfers.

The most visible form of intergenerational transfer is bequests. The importance of bequests, however, is not evidence for a particular structure of preferences. Individuals who leave bequests may be selfish, but unable to spend down their wealth because of uncertainty over the date of death [Davies (1981)] or because of precautionary savings for health risks ('Health Expenditures and Precautionary Savings', in this volume); they may leave bequests as part of risk pooling arrangements with their children ('The

²This argument is made under the partial equilibrium assumption that the interest rate does not change in response to a change in the capital stock.

Family as an Incomplete Annuities Market', in this volume); they may use bequests as a means of payment for services received from children while they were alive [Bernheim, Schleifer, and Summers (1985)]; they may receive utility from the act of giving a bequest [Blinder (1973)]; they may receive utility from their children's utility [Barro (1974)]; or they may just be satiated in their own consumption.

'The Adequacy of Savings' uses data from the Retirement History Survey to address the question of whether individuals save sufficiently for retirement. For each household, the paper calculates the ratio of the stream of consumption that it could afford when it retired to the stream of consumption that it could have afforded at the beginning of its working life. If this ratio is much less than one – that is, if the household could not afford as high consumption once it retired as it could when working – then it is said to have saved inadequately.

The paper finds that most households do indeed have adequate saving. But the interpretation of this fact is difficult: were it not for the forced saving of Social Security and pensions, a large fraction of the sample would have woefully inadequate savings: one third of the couples examined had assets of less than 10% of their future resources (assets plus social security and pension wealth); the figure rises to two-thirds of couples if one excludes housing wealth.³ Looking across individuals, the paper finds that raising the fraction of lifetime income coming in the form of social security or pensions raises the adequacy of old age savings – providing some evidence that this is not a world in which forward-looking, unconstrained agents undo the effects of variations in the timing of their income streams.

5. Thinking about saving outside the life cycle model

The view that the lifetime is too short a horizon over which to look at saving behavior – which is generally associated with Barro (1974) – receives a good deal of discussion in Kotlikoff's work. He clearly sees it as the main competitor to the life cycle model. An approach to saving that is much further outside the framework that Kotlikoff considers – this time in arguing that the lifetime is too *long* a horizon over which to look – comes from the work of Deaton (1989). Because discount rates are high, consumers are liquidity-constrained, and labor income is stochastic, the optimal policy of Deaton's consumers is to hold a small 'buffer stock' of assets. The size of this buffer

³Similarly, Carroll and Summers (1989) find that over one-third of their sample of tax returns report less than \$100 in interest and dividend income in each year of a six-year panel data set, and that only about one-third of the returns report more than \$100 in all six years. Clearly, most people do not have much financial wealth.

stock should be related to the size of income – it might be, for example, a set number of month's worth of income. But the buffer stock is not a function of the individual's stage in his life cycle.

Carroll and Summers (1989) pile up an impressive array of evidence against any far-forward-looking – that is life cycle or infinite horizon – view of saving. They find, for example, that the consumption profile of an individual is closely matched to the income profile of his occupation, even though time paths of income differ substantially and predictably across occupations. Similarly, looking across countries, the rates of growth of consumption and income are almost identical, even though the rate of income growth is fairly predictable. This should not be the case in a forward-looking model where the intertemporal budget constraint, and not current income, is the determinant of consumption.

Much of the contribution of work such as Deaton's is to dispense with the idea that not saving (or saving very little) is inconsistent with utility maximization. One can get this same result in a life cycle framework without relying on high discount rates by invoking social security and liquidity constraints: a large fraction of the population doesn't save for retirement because of social security, and, to the extent that they would rather consume more now, they can't dissave because they can't borrow against their future benefits. Once economists accept that saving may not be optimal for some people, we may be more comfortable with the observation that a large fraction of the population does not save. The observation that some people *do* save, however, raises an intriguing question: what if the life cycle model is appropriate for some but not all of the population?

The idea that saving is best thought about differently for different groups in the population is not without adherents. Carroll and Summers argue that the distribution of saving is so much more skewed than the distribution of consumption, that economists should consider two models: 'one for the liquidity constrained majority of consumers who save little outside of housing equity and one for the small but wealthy minority who seem to do most of the saving' (p. 30). Similarly, Campbell and Mankiw (1989), examining aggregate time series, argue that consumption can be fruitfully broken into two components: part done by liquidity-constrained or rule-of-thumb consumers who spend all of their current income, and part done by forward-looking consumption smoothers. They estimate that about half of aggregate income goes to each group. Burbridge and Robb (1985), examining Canadian data, find that the postretirement path of assets differs qualitatively between income classes: blue collar workers, though they have a low level of assets, look like they are dissaving in accord with the life cycle model, while white collar workers continue to accumulate assets after retirement.

I would argue that the life cycle or the infinite horizon model, whichever is right, is properly thought of as one of three models that apply to different

wealth classes. My proposed categories are the nonsavers, the forward-looking savers, and the non-forward-looking savers. A large part of the population holds almost no assets, or no assets other than a house. It is hard to see these people as *operative* forward-looking savers – although it is not clear whether under some other regime, such as one in which there were no social security, these people would exhibit more saving.

The very rich also seem poor subjects for the forward-looking model. This may be because it is simply impossible for them to consume their wealth quickly enough, as Kotlikoff suggests, or for some other reason – for example, because they value wealth for some reason other than future consumption. Since they appear so rarely in surveys, are reluctant to answer questions, and are probably far from homogeneous anyway, little is known about this group.⁴

In between these two groups are people who accumulate significant wealth but whose utility could be increased by raising consumption above the level that they choose. This is the group for whom the life cycle or infinite horizon models make sense, and for whom the questions addressed in most of Kotlikoff's book are germane.

The relative importance of these three groups depends on the question being asked: the nonsavers may be the largest in number, but contribute little to aggregate saving. This is not to say that this group can be excluded from our thinking about saving. Some types of policy, such as eliminating social security, might turn nonsavers into operative life cycle savers. And policies that prevent nonsavers from slipping into poverty (social insurance) are the same ones that reduce the saving rate of risk-averse life cyclers.

The degree to which the non-forward-looking savers (the very rich) contribute to aggregate saving is a crucial question. It is important to note that the evidence Kotlikoff presents in this book on the overwhelming importance of bequest wealth is not evidence of the overwhelming importance of the rich: forward-looking savers also receive and leave bequests. And to the extent that many of them own houses even though they hold no financial wealth, even nonsavers often leave bequests.

One challenge to any theory of saving is the peculiar role of housing in the portfolios of many households. For example, Merrill (1984) shows that a significant fraction of the elderly population holds a valuable house and few financial assets, and Hurd (1987) shows that how this housing is treated is important to whether one finds the elderly dissaving as predicted by the life cycle model.⁵ We do not know whether the 'failure' of the elderly to dissave their housing wealth represents an unconstrained optimum. Nor do we know

⁴A good place to start is Avery and Elliehausen (1986), which reports the results from the high income supplement to the Survey of Consumer Finances.

⁵Venti and Wise (1989) find that, in apparent contradiction to the life cycle model, those elderly that do sell their houses are as likely to trade up as to trade down.

to what extent the accumulation of housing equity by younger households is forced by constraints in the financial markets.

6. Theory vs empirics

The above section should not be taken as implying that the forward-looking model is not worth studying: it may describe a majority of the population or a majority of the saving that takes place, if not both. What do we know about forward-looking savers? In the conclusion to his introductory essay, Kotlikoff describes the state of knowledge about saving this way: 'A great deal is known at a theoretical level about savings determinants taken one at a time. Much less is known about the interactions of these determinants, and too little is known at an empirical level about the true causes of savings. As is often the case in economics, theory seems to have advanced well beyond the empirical research.' I agree with this summary, but I find the state of affairs to be more alarming than Kotlikoff seems to.

While theory has advanced to study many of the subtleties of possible determinants of savings, examination of empirical literature here and elsewhere leaves one feeling that we cannot be sure of the veracity of even the simplest theory's assumptions or consequences. Are individuals forward-looking in making their saving and consumption decisions? Are there significant altruistic linkages between generations of families? How important are liquidity constraints in determining current consumption? As Kotlikoff describes, testing the theories that we have will require data that may be very difficult to get.

References

- Abel, Andrew, 1985, Precautionary saving and accidental bequests, *American Economic Review* 75, 777-791.
- Auerbach, Alan and Laurence Kotlikoff, 1987, *Dynamic fiscal policy* (Cambridge University Press, Cambridge).
- Avery, Robert B. and Gregory E. Eliehausen, 1986, Financial characteristics of high-income families, *Federal Reserve Bulletin*, March.
- Barro, Robert J., 1974, Are government bonds net wealth?, *Journal of Political Economy* 82, 1095-1118.
- Bernheim, B.D., A. Shleifer, and L. Summers, 1985, The strategic bequest motive, *Journal of Political Economy* 93, 1045-1076.
- Blinder, Alan S., 1973, A model of inherited wealth, *Quarterly Journal of Economics* 87, 608-626.
- Burbridge, J.B. and A.L. Robb, 1985, Evidence on wealth-age profiles in Canadian cross section data, *Canadian Journal of Economics* 18, 854-875.
- Carroll, Chris and Lawrence Summers, 1989, Consumption growth parallels income growth: Some new evidence, NBER working paper no. 3090, Sept.

- Campbell, John Y. and N. Gregory Mankiw, 1989, Consumption, income, and interest rates: Reinterpreting the time series evidence, NBER working paper no. 2924, April.
- Davies, J., 1981, Uncertain lifetime, consumption and dissaving in retirement, *Journal of Political Economy* 89, 561–577.
- Deaton, Angus, 1989, Saving and liquidity constraints, NBER working paper no. 3196, Dec.
- Friedman, Benjamin and Mark Warshawsky, 1990, The cost of annuities: Implications for saving behavior, *Quarterly Journal of Economics* 105, 135–154.
- Hurd, Michael D., 1987, Savings of the elderly and desired bequests, *American Economic Review* 77, 298–312.
- Merrill, Sally R., 1984, Home equity and the elderly, in: H. Aaron and G. Burtless, eds., *Retirement and economic behavior* (Brookings Institution, Washington, DC).
- Venti, Stephen and David Wise, 1989, Aging, moving, and housing wealth, in: David Wise, ed., *The economics of aging* (University of Chicago Press, Chicago, IL) 9–48.