Reflections on Saving Behavior

Robert M. Solow*

Economics needs no special excuse to study the determinants of saving. It is one of those perennial questions always on the agenda. I remember being taught that the Founding Fathers, among them Hume and Smith, had definite views about saving behavior. (I even think I remember being taught that they held a rather sociological theory—thrifty bourgeois merchants, spendthrift hereditary landlords—and *The Fable of the Bees* contains the same suggestion.)

Nevertheless, it is no accident—as we deep-thinkers say—that a conference on government policies affecting saving should be taking place right now. There has clearly been an upsurge of interest in the subject, in this country and elsewhere. One source of curiosity was the apparent fall in reported household saving rates in the United States after 1970 and more particularly after 1975. That may turn out to have been a nonevent, partly a measurement error, partly a short-run phenomenon, partly a shift to other forms of saving, we are still not sure. But whether it happened or not, it helped to focus attention on the saving rate. A more substantial impulse came from the international comparisons, now refined and analyzed in the valuable OECD work of Sturm and Blades. At a time when the U.S. economy felt itself to be losing out in competition with other countries, especially West Germany and Japan, both in international competitiveness and in general economic performance, it was natural to ask: what do they do that we don't do? Clearly one of the things they do is to save and invest a larger share of aggregate income.

Now, of course, raising the saving rate has become a declared object of national policy. The arguments offered on behalf of the policy are not always cogent; and the particular policy measures proposed are not always effective. But it is easy to see why questions about saving behavior are now of special interest. A question can be of interest without being interesting: think of the somewhat related fuss about imminent "capital shortage" just a few years ago. As I mentioned at the very beginning, however, the study of saving is a hardy perennial. It is so closely connected with other aspects of social and economic structure that the basic questions may never be permanently settled. As the dairy industry used to say about milk, you never outgrow your need for the study of saving behavior.

To the eye of an economic theorist, those large international differences in saving rates are the obvious target for explanation. Reasonable people may differ about the capacity of econometrics to make fine analytical distinctions between parameters. But if theory can contribute anything

---

*Institute Professor of Economics at the Massachusetts Institute of Technology.

to the study of actual saving behavior, it ought to be able to give a meaningful account of the possible sources of gross differences like those we see across countries in the figures of Sturm and Blades.

Here a digression is in order. In making international comparisons, we can choose among the household, private, and national saving rates. A believer in the applicability of the Ricardo-Barro equivalence theorem would choose the national saving rate. If the government budget can neither absorb private saving through deficits nor supplement it through surpluses, because households will take whatever offsetting actions are needed to enforce their own intertemporal plans, then it will be the national saving rate that registers the intentions of the private economy. (There is no further implication that public policy is powerless to influence the saving rate. Nobody doubts that a shift from income taxation to a consumption tax would induce an increase in national saving at a given level of economic activity. Any regulation that affects the private incentive to save could do as much. It is only the aggregate budget process that has no force.) Someone like Penner or me, who believes that the Ricardo-Barro proposition is a poor guide to the way the world actually works, will conclude that a country's national saving rate is in part a political decision. Whether that decision is made sensibly or not, economic theory will have little to say about the public component of international differences. The choice of an object for analysis will be between the household saving rate and the private saving rate. I prefer the aggregate private saving rate, because any wealth-oriented theory of saving will suggest that increments of asset value, wherever they are located, will be of approximately equal relevance to the saver. It may be noticed that this is the sort of reasoning that, carried much further, leads to the Ricardo-Barro proposition. That is as it should be, because one of my objections to the Ricardo-Barro view is precisely that it carries a reasonable idea too far, and asks it to function in inappropriate circumstances.

Well, then, why is the private saving rate in the United States x percent and that in Japan 2x percent? There is one sort of explanation that hardly ever occurs to an economist contemplating those figures, but might be the first thought in the mind of a civilian. The Japanese save more because they are the sort of people who save; they are naturally thriftier than we are. Or perhaps they live in a culture that inculcates and values the habit of thrift and the results of thrift more than ours does. Such a response would not have seemed strange to Alfred Marshall who, if not a Founding Father, is certainly a Great Uncle. When he came to consider saving behavior (in the *Principles*) Marshall wrote: "(T)he causes which control the accumulation of wealth differ widely in different countries and different ages. They are not quite the same among any two races, and perhaps not even among any two social classes in the same race. They depend much on social and religious sanctions . . . (T)he chief motive of saving is family affection."

If that were the correct answer to the question, it would be a matter of some importance for the agenda of this conference. The main reason we are struck by the high Japanese saving rate is because we wonder how the U.S. saving rate could be made to look more like it. If the anthropological explanation were the right one, it would suggest strongly—though it does

not prove—that there is not much public policy could do. However much we might wish to achieve a Japanese saving rate, and even if we wished it for the right reasons, it is idle to suppose that we would or could Nipponize our culture, the socialization of our children, our attitudes toward the old or to the family, or whatever it would take to achieve the goal. (This is only a suggestion and not an air-tight argument because nobody believes that nature and/or nurture determine the saving rate to two decimal places. Even an unthrifty society can be induced to save more by economic incentives. It just seems unlikely that differences as large as those between top and bottom in the international pecking-order could be obliterated unless economic incentives were driven intolerably far; and maybe not even then.)

Having raised that question, I want to sketch an economic theorist's answer to it. Most of us hold to a life-cycle theory of saving in one form or another. As several participants in this conference have argued, the life-cycle model is not an unmitigated success; for example, it has difficulty in accounting for the cross-sectional distribution of wealth by age. Nevertheless, that is the vehicle I shall use. For my purpose, familiarity and general acceptability are all-important. It would do me no good to invent an *ad hoc* theory of saving for checking out the significance of "anthropological" variations in the parameters. It comes easily to mind, for instance, to enter wealth itself in the representative saver's utility function (and no doubt this would help to explain the continued saving of the old); but after having remarked sagely that a stronger preference for wealth would lead one to save more, I would have no well-worn standard with which to compare parametric variations.

Even within the life-cycle context, I shall simplify, though I am not happy about that. In particular I ignore uncertainty (because it is difficult to deal with and has ambiguous effects on saving), liquidity constraints (because it is known that their binding presence encourages saving), and social security (because its effects are the subject of current controversy.)

In the absence of uncertainty, then, saving arises for two reasons. The first is what Harrod called "hump-saving": accumulation and decumulation of assets arising because the representative household wants a lifetime consumption pattern that is smoother than its lifetime profile of earnings. Saving for retirement is the most important factor here. The hump of assets is built up during the working life and run down during retirement. It is well understood that this mechanism causes the saving rate to depend on the age distribution of the population. A rapidly growing population will have relatively fewer households of retirement age and relatively more still working, and therefore generally a higher saving rate. (The burden of supporting minor children must also be factored into such calculations.) The second reason for saving is the accumulation of a bequest to be passed on to the next generation. (The existence of uncertainty about income and costly contingencies generates a precautionary motive for saving as well.)

Sociological and cultural influences on saving behavior enter an economic model like this primarily through parameters describing tastes. In the life-cycle framework, there are three points at which tastes must be specified: (1) a time-preference or utility-discount rate, (2) a weight to be

attached to the bequest motive, and (3) the choice of an instantaneous util-ity-of-consumption function, especially its concavity, which will affect the desired degree of consumption-smoothing. It is worth pointing out that the major analytical discussion in the literature of international comparisons from the life-cycle point of view—Franco Modigliani's 1970 article—sim-ply dismisses the possible significance of intercountry differences in taste with the remark that the parameters of the saving function do not seem to be very responsive to such forces within the relevant range of variation. Modigliani concludes that ". . . all the evidence supports both qualitatively and quantitatively the role of the two principal variables suggested by the life cycle model, productivity growth of income and the age structure of the adult population. Furthermore, these variables appear to account for two-thirds to four-fifths of the inter-country variance in the saving ratio." That would leave some room for policy, but not a lot, because it is far-fetched to imagine the age distribution as a tool or object of policy, and the rate of productivity growth is not easily controllable either, in practice.

Until quite recently there has been little discussion of even something as "obvious" as the bequest motive as a source of intercountry differences in saving rates. I have the (casual) impression that this neglect was more or less accidental. In the early days of life-cycle theory, it was natural for Modigliani and others to assume away the existence of bequests in the search for the simplest formulation of the theory that would highlight its most novel feature, the emphasis on hump-saving. Then a well-known pa-per by James Tobin used carefully constructed numerical examples to show that hump-saving alone could generate substantial net saving in a growing population. Tobin concluded that ". . . it seems quite possible that life cycle saving can account for the United States capital stock." The combination of convenience and parsimony tended to favor playing down the bequest motive.

Lately, however, the opposite conclusion seems to have gained force. Numerical calculations by Betsy White imply that pure life-cycle saving by itself can not generate a realistically high volume of saving. Unfortunately such numerical experiments are not very transparent. Söderström com-ments: "The reason why White and (Tobin) reach such different conclusions is not quite clear. Since their models are similar in general design, the rea-son has to be sought in differences of details, and those are numerous . . . (O)ne can only say that details are very important." That is bad news for the theorist. Nevertheless, White's results at least suggest that it might make sense to supplement hump-saving with the desire to accumulate a bequest. Söderström's own contribution is entirely consistent with this suggestion. He too does numerical exercises with a model in which there are no planned bequests. But there is uncertainty about time of death, and so some individuals die unexpectedly soon, owning positive net worth which is passed on to survivors. These transfers (also unexpected) will be saved, at least initially, and show up in measured aggregates. In Söderström's formu-lation of the life-cycle model, these unplanned bequests turn out to be very important. The model can generate saving rates as high as those observed (in Sweden, to which the model is calibrated) and even higher.

Finally, Kotlikoff and Summers proceed differently, though still within the framework of the life-cycle theory of saving. Instead of making hypothetical calculations to show what the theory would predict for reasonable parametric specifications, they look at profiles of earnings and consumption by age to estimate the stock of life-cycle wealth directly. They find that life-cycle wealth accounts for only a small fraction of aggregate private wealth in the United States. They say: "The evidence presented in this paper rules out life-cycle hump saving as the major determinant of capital accumulation in the U.S. economy. Longitudinal age-earnings and age-consumption profiles do not exhibit the kinds of shapes needed to generate a large amount of life-cycle wealth accumulation . . . Intergenerational transfers appear to be the major element determining wealth accumulation in the United States."

My tentative conviction is that this view of the matter is essentially right. It is reinforced by general qualitative considerations. The natural temporal habitat of a theory emphasizing the life cycle ought to be periods of 50–100 years. Over historical time, the impression is inescapable that successive generations start (economic) life with larger per capita endowments of tangible wealth than their predecessors. Since the distribution of wealth is known to be highly concentrated, this description need only apply to the relatively small group of families owning among them a large share of the capital stock. But then it becomes important to understand how the bequest motive influences the saving rate.

I want to report one primitive experiment along this line. It is too crude to provide more than a hint that there is something worth exploring, but I have limited myself to what could be done with pencil and paper and a small hand calculator. The version of life-cycle theory I adopt is essentially that described by M. Yaari.[1]

The representative individual is born (at age 20, say), works for 45 years, lives in retirement for 10 more years, and dies. The population is constant and so, in this one-person-shay world, the age distribution is uniform. An individual born at time zero starts with an inheritance that I set equal to some multiple of the annual wage at time zero. (Later I put that multiple equal to one, so the representative person inherits wealth whose present value discounted back to age zero (calendar age 20) is one year's pay at the wage ruling at time zero.) The annual wage is independent of *age* but grows exponentially through time at $100g = 2$ percent a year. So everybody of working age at time t earns $y_0 exp(.02t)$. At birth, the individual disposes of the present value of the rising earnings over a 45-year work-

ing life, plus the inheritance already described. The (real) interest rate is constant at 3 percent a year.

The same individual, at birth, plans consumption over a 55-year life span. S/he maximizes

$$\int_0^{55} e^{-jt} u(c)dt + kv(b)$$

where j is a rate of time preference, $u(.)$ is the instantaneous utility-of-consumption function, and $kv(b)$ is the utility of contemplating a bequest of amount b to one's descendant. I imagine the bequest to be paid out at death. The number k is a parameter carrying the strength of the bequest-motive. I set $j = .01$ and have not bothered to try other values. We know about the qualitative influence of time preference on the propensity to save. At $j = .01$, a person at age zero discounts utility at age 55 by about one-half. Some economists have experimented with rates of time preference as high as $j = .04$. In that case, the discount factor over a 55-year life span is one-tenth; from the vantage point of age 38, I can *feel* that must be wrong. As usual, I have taken $u(c) = c^{1-h}/(1-h)$ and experimented with a few values of h in the range from $h=1$ to $h=2$. I have also taken $v(.)$ to be the same function as $u(.)$, even to the same value of the elasticity parameter. It is simply a great arithmetical simplification to do so, and there is nothing much against it. I will come back to the bequest-parameter k later on.

The choice of a consumption profile is subject to a constraint that I need not write down in detail. It merely requires that the present value of the inheritance plus the present value of earnings over the working life equal the present value of the bequest plus the present value of consumption over the life span. All this discounting is done at the interest rate $i(=.03)$, so there is a tacit assumption that the capital market is perfect. Many others have studied the effect of a social security scheme in this context, so I ignore that question.

The rest is routine. Optimal consumption at age a is proportional to $exp((i-j)a/h)$. Since I have the interest rate (.03) bigger than the utility discount rate (.01), optimal consumption grows throughout the lifetime. The return on savings is big enough to overcome normal impatience. If $h=1$ (the logarithmic utility function), consumption grows at 2 percent a year. If $h=2$, so that the marginal utility of consumption falls considerably faster, consumption grows at only 1 percent a year over the lifetime; greater smoothing is achieved.

The *level* of the consumption profile and the size of the optimal bequest are determined together to meet two conditions. Consumption at any age could always be made a little bit larger or a little bit smaller; with due allowance for interest earnings, any such variation in saving can be translated through time into a corresponding change in the bequest. Optimality requires the obvious utility-balancing at the margin. In addition,

of course, the consumption level and the bequest are tied together through the lifetime budget constraint.

The first result I want to report is about the sensitivity of the saving rate to variations in the bequest parameter (called k). The model can be calibrated to give results that are not outlandish. For instance, with the specifications already made (time preference rate $= .01$, productivity growth rate $= .02$, interest rate $= .03$, working life $= 45$ years, lifetime (from age 20) $= 55$ years, inheritance at age zero $=$ one year's wage at that time), the choice $h = 1$, $k = 0.7$ leads to the following results: the optimal bequest is (approximately) equal to one year's wage at the time of bequest, and an aggregate saving rate of about 12 percent. The optimizer thus leaves a bequest equal to $exp(55(.02)) = exp(1.1) = 3$ times the inheritance s/he had received 55 years earlier, but that just allows for rising incomes. In calculating the saving rate, national income is defined as aggregate wage income plus interest at 3 percent on the aggregate wealth of the population summed across the uniform age distribution. By the way, with these parameters, the wage share in aggregate income is about 85 percent, and the wealth-income ratio is about 5:1.

Now suppose the bequest parameter is changed to $k = 0.8$, with the rest of the specification as before. The size of the optimal bequest rises by about 15 percent. But the aggregate saving rate does not change at all to two decimal places. If k is reduced to 0.5, the optimal bequest is 30 percent lower than with $k = 0.7$, but again the saving rate moves imperceptibly.

Obviously I should spend my spare time replicating these calculations for other parameter-sets, for a growing population, for alternative initial inheritances, etc. For now, I interpret the model as saying that modest changes in each generation's concern for its heirs will have very little effect on the aggregate saving rate. Since more than modest changes in such attitudes are not likely to be achievable, there is not much to be gained in that direction. The intuitive reason why modest changes in the bequest motive have so little effect on the saving rate must be something like the following. If the bequest is the order of magnitude of one year's income, then the intention to accumulate a somewhat larger estate will not call for any large increase in saving in any given year. The added saving effort will be spread over a lifetime according to the standard marginal equivalences. The effect on the saving rate will be further moderated by the extra interest income earned *en route*. It seems likely that this generalization will be approximately true even if the target bequest is a bit bigger than one year's earnings so long as it constitutes a number of years' earnings, that is, small compared with the working lifetime.

The implication for international comparisons is trickier. It is clear from the mathematics that a sufficiently large value of k can drive consumption toward zero. So it is possible in principle that Japanese save more than Americans primarily because they just get much more satisfaction from providing for their heirs. Somehow I doubt that; but it would take comparative data on inheritances and incomes to check it out.

The second preference-indicator is h, the concavity-parameter. Here the story is rather different. I go back to $k = 0.7$. With logarithmic utility, $h = 1$, the bequest keeps pace with earnings, and the aggregate saving rate is 12 percent. Let $h = 1\frac{1}{2}$, so the representative saver experiences more sharply diminishing marginal utility of consumption and thus is more risk-averse, more desirous of smoothing consumption over the lifetime. Someone with $h = 1\frac{1}{2}$ leaves a bequest about a quarter smaller than someone with $h = 1$, but such a population's saving rate falls only to 11 percent. If, however, we set $h = 1.75$, the bequest falls another 6 percent below the level at $h = 1$ and the saving rate drops to $7\frac{1}{2}$ percent; and at $h = 2$ (a value recommended by Ragnar Frisch, I think) the saving rate falls drastically to 4 percent. Evidently the effects of this parameter are strongly nonlinear, and the impact on the saving rate can be dramatic.

Remember that these calculations automatically equate the elasticity parameters in the utility functions for bequests and for own consumption. That was done for mere arithmetical convenience, though it is perhaps not bad to assume that tastes which saturate more quickly with respect to consumption will do so also for bequests. Anyway, it is pretty clear that the main effect of sharper concavity is to reduce the desired rate of growth of consumption over the lifetime, to favor the early low-income years. It seems to me entirely possible that the sort of difference in tastes captured by variations in h could dominate international comparisons.

That would be a pessimistic conclusion from the policy standpoint. There are some puzzles, however. It goes against deeply entrenched clichés to conclude that Germans and Japanese save more than Americans because they are insatiable gamblers, while we are unambitious conservatives who soon tire of the pleasures of high consumption. But of course the difference, if there is one, need not refer to personality type. It could reflect institutional differences that favor or disfavor the conversion of a rising income stream into a more nearly level consumption profile. That effect could be modified by policy, though it is hard to imagine how a democratic government would justify depriving its citizens of borrowing opportunities they have customarily enjoyed.

Another cliché, that Americans are more oriented toward the short run than others, that we "want it now," does suggest itself. That sounds more like a statement about time preference, however. I commented earlier that small values of the time-preference parameter (in the neighborhood of .01) sound most plausible. A little experimentation shows that minor variations around so small a rate of discount do not have dramatic effects on the saving rate. The possibility remains that there may be very large intercultural differences in time preference. They would, of course, strongly affect the saving rate in a well-understood way. It is not clear that economic policy could, or should, do anything to increase a private saving rate that is low because citizens do not care much about the future. If there is a social rate of time preference that is systematically smaller than the private rate, for the sorts of reasons that have been discussed in the literature, it would seem more straightforward to give effect to the social interest through the government's contribution to the national saving rate.

My conclusions from these "anthropological" experiments are on the whole pessimistic. Small changes in the bequest-motive and in the rate of time preference have only small effects on the saving rate. Drastic changes would no doubt have large effects, but there is no opening for economic policy in that observation. No democratic government could or should try deliberately to generate big changes in the preferences of its citizens, especially not when the parameters in question may have deep cultural roots and are not obviously self-destructive.

The case of the concavity parameter is slightly different. It is far from clear what constitutes a "small" difference in the elasticity of the marginal utility of income, but one has the impression that the saving rate is fairly sensitive to changes in the speed with which marginal utility diminishes. It is tempting to think of this parameter as measuring risk-aversion; if it is "too much risk-aversion" that makes the saving rate "too small" then, even if public policy can do nothing about preferences, it can certainly do something about the degree of risk to which private savers are exposed. But there is no uncertainty in the simple model I have used as a trial horse. We know from other models that do admit uncertainty that there are offsetting effects to be dealt with. Strong risk aversion (rapidly diminishing marginal utility of income or wealth) favors saving to protect the saver against painful unfavorable contingencies, and also works against saving because the added interest income is not much valued. Under certainty only the latter effect operates; a high value of h goes against saving in order to achieve a flatter lifetime consumption profile. Somehow I doubt that governments will have much luck in preaching either the Puritan ethic or the pleasures of the pot of gold at the end of the rainbow.

The life-cycle model has something to say, in principle, about the effectiveness of interest rate variations in stimulating saving. As Marshall knew, of course, there are offsetting effects; bequests, for instance, are cheaper to accumulate but it takes less saving to accumulate them. The net effect can go either way. In the model used here, it turns out that a rise in the real rate of interest from 3 to 4 percent does indeed increase the saving rate: from 12 to a bit under 15 percent in the base case. Tobin does not address that question directly in his simulations with a life-cycle model, but his calculations seem to imply a similar favorable effect of higher interest rates on saving. White's calculations with an apparently similar model produce the reverse effect: a higher interest rate is accompanied by a lower saving rate. This finding points up Söderström's discouraging remark that the details seem to matter in these models. It might be worth someone's effort to discover just why these two experiments yield contradictory answers to important questions, but it is hard to imagine who would go to the trouble.

When theory suggests offsetting forces, one says we are faced by an empirical question, but in some ways that is even more discouraging. Over the years, we had come to accept a piece of folk-econometrics: no one had succeeded in finding a reliable interest rate effect on saving, and so macroeconomics proceeded on the presumption that the saving rate was effec-

tively independent of the interest rate. But now first Boskin and then Summers claimed to have found a significant and substantial positive effect in U.S. time series. In principle, why not? The passage of time adds new observations, more sophisticated statistical techniques, and clever new formulations. Still more recently, however, Friend and Hasbrouck have reviewed the new evidence and concluded that it will not stand up. It is too sensitively dependent on particular definitions, particular sample periods, and particular formulations. The only robust result is the old folk result.

The sad part of this story is the suggestion that applied econometrics could become a forensic subject. One fears that econometric testimony about the effect of interest rate changes on saving—or the effect of unfunded social security wealth on private saving, to take another prominent and relevant example—is on the same footing as psychiatric testimony about John Hinckley's sanity. That is to say, it provides moral support for one side or another, but no one would confuse it with scientific evidence. In this state of affairs, it seems plain that robustness is the prime econometric virtue. On that basis, there is no solid intellectual foundation for the notion that feasible variations in the after-tax return to private saving are a good way to increase the share of private income saved.

Economic theory could have something to say about a logically prior question: is there any good reason to desire a higher private saving rate than there actually is? The natural place to look for an answer is in the theory of optimal capital accumulation, with the caution that so abstract a theory can not be asked for more than order-of-magnitude indications of the socially optimal saving rate. To avoid complications, I shall limit myself to steady-state calculations.

Right at the start, there is a subtle distinction to be made. The literature contains two versions of the quantity "society" should be thought to be maximizing. One version suggests

$$\int_{0}^{\infty} e^{-at} c^{1-h}/(1-h)\,dt.$$

The notion is pretty much as before: $c^{1-h}/(1-h)$ is the (social) utility of a per capita rate of consumption equal to c, and a is a rate of (social) time preference. This discounted social utility is to be maximized subject to initial conditions, a given time path for the labor force, and a technology that converts accumulated capital and available labor into an output that can be used either for current consumption or for capital accumulation. Under generally acceptable technical assumptions, it is traditional to show that the optimal policy drives the saving rate to a steady-state value

$$s^* = \frac{gb}{a+n+hf}$$

where b is the (Cobb-Douglas) elasticity of output with respect to capital input, f is the rate of (labor-augmenting) technical progress, n is the rate of growth of the labor force, and $g = n + f$ is the "natural rate of growth" of the economy.

I happen to prefer this formulation. The literature also suggests an alternative, in which the integrand above is multiplied by e^{nt}. The idea is that society "credits" itself at each instant not with the discounted social utility of consumption per head but with that quantity multiplied by the number of people enjoying the representative utility level at that instant. There is no point in arguing the pros and cons here; obviously this alternative formulation leads to a larger optimal steady-state saving rate if the population is growing, because it pays to shift consumption to the future, when there will be more people. This higher optimal saving rate can easily be seen to be

$$s^{**} = \frac{gb}{a + hf}.$$

To get a feel for the magnitudes, suppose we put $n = .01$ and $f = .02$ (although this may be an optimistic figure after so many years of stagnating productivity). Then $g = .03$. I continue with $a = .01$, as if the social rate of time preference were the same as the private discount rate used in the life-cycle calculations. I owe to Peter Sturm a reminder that the use of these formulae as bench marks calls for explicit attention to the difference between gross and net saving and gross and net income. In theory, optimal saving rates are net concepts. In practice that means being careful about the value assigned to the elasticity b. (I would add that the breadth of the capital concept must also affect b.) In the United States, capital consumption allowances in the national accounts run about 10 percent of GNP. That suggests a value of b between, say, .15 and .20.

Then s^* ranges from 11 to 15 percent if $h = 1$, and from 8 to 10 percent if $h = 2$. Setting the social rate of time preference equal to zero, as Frank Ramsey thought proper, would increase s^* to 15–20 percent with $h = 1$ and 9–12 percent if $h = 2$. Going back to $a = .01$ but setting $f = .01$ leaves s^* at 15–20 percent if $h = 1$ and makes it 11–15 percent if $h = 2$. The alternative formulation leads to larger values of s^{**}, going as high as $22\frac{1}{2}$ percent at its maximum.

What should a reasonable person make of this? The numbers are not to be taken literally, of course. But I do not think they are totally irrelevant. Policy talk sometimes seems to take it as axiomatic that it would be a Good Thing to promote a higher saving rate. Even an excessively formal model like this one has the merit of forcing one to provide reasons. The classical reason for Growthmanship still holds: if private savers discount the future more drastically than "society" ought, then private saving will fall short of the social optimum. One could make the same argument in terms of a difference in the relevant risk premia for private decisions and social decisions. I notice with regret that the slippery concavity parameter h turns out to be important here too, but I don't know that one can make much of that, for reasons already mentioned.

Do we actually save too little? Presumably it is the national saving rate that should be compared with s^* (or s^{**}). Presumably also at least some of the Sturm-Blades adjustments should be made to the current figures from the national accounts; logic would seem to call for the inclusion of consumer durables and at least public nonmilitary durables in the total. In those terms, the gross national saving rate in the United States is at least 24 percent, and shows no downward trend, according to Table 13 in Sturm and Blades. In net terms, the adjusted saving rate must be close to 15 percent. That is in the ballpark bounded by the theory of optimal capital accumulation, and could even be on the high side. Even without consumer durables, their figure is 18.6 percent. Robert Eisner rightly calls attention to the fact that these are all steady-state calculations. If other countries are moving up to their steady-state capital stocks, while the United States is more nearly there, a further reason for the observed international difference emerges. Leaving aside the arguments about divergence of private and social costs and benefits, the possibility exists that it doesn't matter much that we don't know how to fix the saving rate, because it ain't broke.

There is, after all, something fishy about the current enthusiasm for saving incentives. In the first place, as I have just argued, no very good reasons have been proposed for believing that the national saving rate in the United States is too low. (I do not count merely pointing at Japan as a reason, although I would be happier if we understood the Japanese economy better.) Nor do the data of Sturm and Blades suggest that the national saving rate has been falling. Taking a longer period of time would almost certainly reinforce that finding.

Secondly, as Jump argues convincingly for Canada, the incentives that have been proposed and enacted for increasing the household saving rate are uncertain, maybe even perverse, in their effects. In this country, no one argues that the All-Savers Certificates, for instance, generated any measurable *ceteris paribus* rise in personal saving. If it were desirable to engineer an increase in the household saving rate, there are more effective ways: a shift from the taxation of income to the taxation of consumption is one obvious possibility, though I hasten to add that so drastic a change would need to be considered from many points of view, especially since the need for more personal saving is not firmly established. The particular incentives for saving that have been proposed—in Canada too, apparently—are so uniformly in danger of turning into mere lump-sum transfers to the well-off, that a person with even an ounce of healthy cynicism is bound to wonder if that were not their primary purpose.

Third, whatever the long-run need for national saving may be, there is no intelligible case to be made that the volume of plant and equipment spending in the United States is now or in the near future limited by an inadequate willingness to save. At a time when capacity utilization rates are below 70 percent and not visibly going up, it is hard to believe that an investment boom is waiting to be touched off by incentives for personal saving. (I noticed that Secretary Regan is trying to direct the public's attention away from the prime rate, because he thinks that it gives the impression

that the level of interest rates is higher than it actually is. He said something about our need to get people into the banks and borrowing. Are we all Keynesians again?) As several papers in this conference have noted, it would make far more sense to be thinking about sweetening the inducement to invest than about raising the propensity to save, and Auerbach's paper shows how it can be done.

Fourth, this whole line of argument suggests the wisdom of a policy direction that will be nostalgically familiar to at least two Old Growthmen in this group: a shift toward easier credit and tighter fiscal policy, with the first looking after the inducement to invest and the second providing the saving through the public budget. It goes without saying that Old Growthmen know full well that this is a policy that makes sense only when the economy is operating near its normal capacity to produce.

Which brings me to my fifth and last point. There is an incomparably larger source of personal saving waiting to be tapped than any fiddling with the taxability of interest could generate. If I may give it a name, how about the multiplier process? The GNP gap can hardly be a lot less than 9 percent of current output. I do not know what the best estimate of the marginal national propensity to save is, but if the short-run multiplier is a little less than two, the national marginal propensity to save—counting in retained earnings and induced tax revenues—must be close to one-half. So a mere closing of the gap would add 4 percent of GNP to the national saving rate. I understand the dangers of overheating an economy not yet out of the inflationary woods; and I realize that any initial expansionary impulse might come in part from a step reduction in public saving. The point of this exercise is only to underline the wisdom of Penner's willingness to trade stimulus later for stimulus now—which I would make more overt—and to remind you "where the saving will come from."

REFERENCES

Boskin, M. "Taxation, Saving and the Rate of Interest," *Journal of Political Economy,* Vol. 86 (1978)

Friend, I. and J. Hasbrouck. "Saving and After-Tax Rates of Return," Mimeo (1982).

Kotlikoff, L. and L. Summers, "The Role of Intergenerational Transfers in Aggregate Capital Accumulation," *Journal of Political Economy,* Vol. 89 (1981) 706–732.

Modigliani, F. "The Life Cycle Hypothesis of Saving and Inter-country Differences in the Saving Rate," in W. A. Eltis, M. F. G. Scott, J. N. Wolfe, eds. *Induction, Growth and Trade,* (Essays in honor of Sir Roy Harrod) Oxford: Clarendon Press, 1970, 197–225.

Söderström, L. "The Life Cycle Hypothesis and Aggregate Household Saving," *American Economic Review,* Vol. 72 (1982) 590–596.

Summers, L. "Capital Taxation and Accumulation in a Life-Cycle Model," *American Economic Review,* Vol. 71 (1981).

Tobin, J. "Life Cycle Saving and Balanced Growth," in Fellner et al., eds., *Ten Economic Studies in the Tradition of Irving Fisher,* New York: Wiley, 1967, 231–256.

White, B. "Empirical Tests of the Life Cycle Hypothesis," *American Economic Review,* Vol. 68 (1978) 547–560.

Yaari, M. "On the Consumer's Lifetime Allocation Process," *International Economic Review,* Vol. 5 (1964) 304–317.

Discussion

Henry C. Wallich*

When I read Bob Solow's elegant paper, I thought I had a few things to say. After then looking at all the other papers, I realize that everything has been said before. I can only make a few remarks from, as it were, an intuitive perspective rather than with theoretical rigor or econometric analysis. I fear that Bob is correct in suggesting that this discussion of whether to stimulate saving has some of the aspects of the discussion of John Hinckley's sanity. It is always pro domo. No matter what you hear about what kind of tax change is needed, it always ends up with a tax benefit for the speaker or his constituency. I am just as suspect here as anybody else.

I share the perplexity as to the need for accelerated growth. I used to feel quite strongly about that need, mostly on the grounds of international power relationships and also in order to ward off domestic pressure for redistribution. I feel a little less hopeful now that those things are going to be resolved by accelerated growth, but I think growth is still a desirable objective. However, I also feel uncertain whether accelerated investment is the way to get accelerated growth. Bob Solow was one of the first to cast doubt upon that seemingly obvious proposition. One proposition that I think is reasonable is that if we want more investment, then we need more savings. To fiddle with investment makes sense only if that produces more saving. The only way in which encouraging investment will raise saving is if by raising investment, we first raise interest rates and if saving then responds positively to that higher interest rate. As Bob has pointed out, on a folkloric basis that is a mute subject. We do not know if there is an effect or not.

As far as Bob's parting shot is concerned to the effect that given present-day excess capacity and unemployment, there is really no need to worry about saving, I would say yes to that in a short-run view. But when one fiddles around with anthropological parameters and thinks that maybe one can change various propensities of the population, one does have to allow some time. Over that time, the Federal Reserve and the Congress may get us back to a higher rate of capacity utilization.

Let me look at some of the motives for saving that appear in Solow's paper, beginning with the life cycle hypothesis. Somehow it seems to me that in this particular area, economists are dealing with motivations as if they were looking at an alien species; they are not looking at human beings but at an ant hill and are trying to explore why they run in one direction or another. It seems to me that the life cycle hypothesis violates a very basic instinct—the acquisitive instinct. People do not like to decumulate. Once they have got something, they like to hold on to it. They are aware that they

*Member, Board of Governors of the Federal Reserve System.

cannot take it with them, and they are also aware that if they could, it would only melt there. Nevertheless, I think there is a profound reluctance to use up one's savings. You see that in the way people have organized their savings. If we wanted to implement the life-cycle hypothesis in our personal lives, then presumably we would all be buying annuities. Now we are, of course, buying some annuities as well as relying on Social Security. We rely quite heavily on pension funds and some life insurance. These usually take the form of an annuity. But the great bulk of savings, I think, is still in discretionary rather than contractual form. People could have changed that by relying more heavily on pension funds and on life insurance but they have not. They invest in deposits, they invest in money market instruments, they invest in securities and real estate. The Federal Reserve flow of funds tells us that these are the bulk of investment.

Now look at the bequest motive. Everybody has a desire to leave his children a good education. That is one bequest everybody makes who can. They like to make this bequest early in life. Perhaps that is another reason for including investment in human beings in total savings, with education as part of total investment expenditure.

Other than that, the bequests that Bob analyzes are really very small ones. Two or so years' income is not very much, and it is not only too little, it is also too late. By the time a person's life expectancy expires at 75 or so, his children are not very far from retirement themselves and they do not need the bequest. Moreover, they are probably making a great deal more than daddy, at least if their daddy is in the academic line, and so he really is not doing a great deal for them. There may be some sense of social obligation here, and being near Newburyport puts me in mind of the late John P. Marquand and *The Late George Apley.* You may recall in that book that there is a scene where George's father tells him how to use his money. You take your income, the father says to George, and divide it in two equal parts. One part you set aside for saving. He does not say for bequests, he just says "save it." The other half you again divide into two equal parts, one of which you give to charity, and the other half, which is one-fourth of the total income, you spend for your living needs. Now that was a good bourgeois attitude, just like the bourgeois attitude that a man should not retire until he has enough to maintain his standard of living in retirement and still save a little. As you know, the data we have on our aged seem to say that although they may not be maintaining their standard of living, nevertheless, they are still saving. This acquisitive instinct, I think, is at work rather than a bequest motive.

Last, I turn to the precautionary motive that Bob mentioned. The paper first eliminates uncertainty and later admits it. By eliminating uncertainty, Solow arrives at a very peculiar proposition, namely that risk aversion reduces saving, and that is quite true. If there is no uncertainty, then the concavity parameter has only the effect of making people want to have a flat income stream throughout their lives rather than to accumulate capital. Therefore, they do not react to interest rates and they save little. But as soon as you give up the certainty postulate and admit uncertainty, there is a very powerful precautionary motive for saving. The great risk of the

saver, of course, is to outrun his original life expectancy. If he does not want to invest in an annuity that would eliminate that risk, he can avoid it by planning on living to 100 or 110. We know, of course, that no one ever reaches his full life expectancy on the day of his death, but looking ahead there is a great deal of exposure to be covered if one wants to be completely sure. We see the strength of the precautionary motive in the differential saving behavior in particular social groups. Farmers save more because they are exposed to risk, the self-employed save more because they are exposed to risk. That seems to be a very powerful factor.

As for the determinants that take the form of interest rate and, I would note, inflation, I have nothing to add to the findings. We have a study at the Board by Steindel which does find a positive effect of interest on saving. Bob knows that study, and I am not telling him anything new. I could give you some purely anecdotal evidence. We get complaints from automobile dealers who report people come into their showroom and say they have $5,000 or $10,000. They say they could buy a car for cash but they are getting 15 percent on their money market mutual fund and they hate to give that up. That seems to be an influence of the interest rate on saving even though the person probably has money illusion because after taxes he is probably still getting the negative real rate.

As for the effect of inflation on saving, that seems to pose a problem similar to that of the rate of interest on saving. Is it positive? Is it negative? Inflation influences the real rate of return. What is the real rate of return after tax? That we have to think about. We do not know which way that pushes the saver. We do know that the borrower borrows more as inflation reduces his real rate after tax.

But far more important, of course, is the effect of inflation on accumulated past saving. If we were talking here about what I would really like to discuss, we would not talk about how to increase saving, but how to make it possible for people to keep the savings they already have. I know, of course, that for everybody who loses from inflation there is somebody else who gains, including these days the federal government which behaves exactly as if its dissaving were responsive to inflation. Its debt is reduced and it responds by borrowing more. In any event, the impact of inflation on past saving, through its redistributive effect, seems to me to be far more serious than the impact on new saving whatever that would be. It creates a degree of uncertainty in people's lives—the difficulty of knowing how to put your children through college, the difficulty of knowing how to provide for old age. It is very curious that, in the midst of this uncertainty, there have not been greater changes in the savings ratio of the United States than we have observed.

I suppose you were made aware yesterday there seemed to be a tremendous drop in the savings ratio in the United States in the second half of the 1970s. Excellent papers were written about it. I have with me a paper here from the Federal Reserve Bulletin dated August 1980 that analyzes this phenomenon. Unfortunately, in January 1981 and July 1982 all these shortfalls were largely revised away. Now we have got to find a new set of

reasons for a different set of phenomena. The savings ratio has indeed been remarkably constant and this is in the face of a great deal of uncertainty on the part of people who do the saving. But these people have one other way out and that is to rely increasingly on the federal government. Under conditions of inflation, you cannot rely on what you may get from your insurance company. You cannot rely on the money you put in the savings bank, or government bonds, because it may be inflated away. The only party that can give you certainty is the federal government with an indexed pension and Social Security. That is a very troublesome situation. It is surprising that we have not developed indexed bonds. I cannot believe that the risk would be all that high even for private issuers. If firms can deal with 10 percent wage increases, why can't they deal with an increase in interest costs? I think the real reasons why private borrowers have not developed indexed bonds is that interest is tax deductible. Presumably the inflation premium in an indexed bond would not be deductible.

Let me turn to some of the measures we might take other than to manipulate the anthropological parameters. First, it seems to me that something could be done about the income tax treatment of inflation with respect to holding gains. I realize that as soon as someone says do not tax inflationary holding gains, somebody else will mention the tax deductibility of interest. These two things have to be weighed against each other. Something clearly should be done about the capital gains tax. The British have done something very interesting. All inflation gains after a certain date are nontaxable; that has enabled them to issue an indexed bond and to make the inflation premium in that bond nontaxable because it is identical with the rate of inflation. Today, the capital gains tax for people who acquired their assets long ago is really a tax on capital and not on gains. Something could be done about the estate tax if one were really minded to protect saving. The estate tax is paid almost wholly out of saving. It presumably ought to influence people's desire to make bequests but apparently all it does is influence their desire to make a tax-exempt bequest to Yale, Harvard, or MIT rather than leave it all to their children. There is the possibility, that Bob discussed, of a consumption tax. One could certainly shift the balance of saving and consumption in that way. Finally, a budget surplus to my mind would add to saving. I cannot believe, and I do not see, that people respond to this large government deficit by increasing their own saving. At least in the short run, there is a strong dissaving element in the posture of the federal budget.