THE EFFECT OF INTEREST-RATE CHANGES ON HOUSEHOLD SAVING AND CONSUMPTION: A SURVEY

Douglas W. Elmendorf
Federal Reserve Board

June 1996

I have received helpful comments from Bob Dennis, Karen Dynan, Bill Gale, Joyce Manchester, Greg Mankiw, Ben Page, Bill Randolph, Frank Russek, Kent Smetters, and John Sturrock. I am grateful for research assistance from Sarah Levendusky, John Romley, and Jennifer Wolfson. Much of this paper was written while I was a Principal Analyst at the Congressional Budget Office. The views expressed here are my own, and not necessarily those of the Congressional Budget Office or the Federal Reserve Board or its staff.

Abstract

Direct estimates of the interest elasticity of saving suffer from several serious problems. As an alternative, this survey uses an indirect approach that combines models of individual behavior with estimates of certain features of individuals' preferences. The paper examines the effect of interest-rate changes on the consumption and saving of people who follow the lifecycle model, who plan to leave bequests, who save to reach a fixed target, and who have short planning horizons.

It is not possible to provide a precise estimate of the interest elasticity of saving with any confidence. Nevertheless, the models that likely describe the behavior of the people who account for most of aggregate saving imply positive elasticities, so the aggregate interest elasticity of saving is probably positive.

1. Introduction

Understanding the response of personal saving to changes in interest rates is central to many issues in economic policy. For example, a reduction in the budget deficit would probably cause interest rates to decline. If personal saving declined as a result, the overall increase in national saving would be less than the reduction in the budget deficit. Alternatively, contractionary monetary policy generally causes interest rates to rise. If personal saving increases as a result, the corresponding fall in consumer spending helps to slow the economy. As a final example, changes in the tax code can raise or lower the net-of-tax return to saving. The effect of these changes on the amount of saving may play an important role in tax policymaking.

The interest elasticity of saving is defined as the percent change in saving that results from a one-percent change in the interest rate. There is disagreement among economists about both the sign and magnitude of this elasticity, as existing theory and empirical evidence do not appear to offer any clear conclusions. The goal of this paper is to review economists' knowledge about this critical aspect of people's behavior.

Economists' standard model of consumer behavior is the lifecycle model, which assumes that people determine their consumption and saving at each point in their lives by looking forward to their future income and desires, rather than considering only their current income and desired spending. Basic economic courses use a stylized version of this model to show that the interest elasticity of saving can be decomposed into a "substitution" effect and an "income" effect, which work in opposite directions. Although this characterization of the lifecycle model appears in most casual writing on the topic, it is appropriate only for individuals in very specific situations. In particular, this characterization is incomplete because it ignores the way in which interest rate changes induce revaluations of existing wealth that affect saving. This additional "wealth" effect contributes positively to the interest elasticity of saving, thus reinforcing the substitution effect, although it does not resolve the theoretical ambiguity about its sign.

Alternative models of consumer behavior lead to different analyses of the interest elasticity of saving. While lifecycle consumers consider all of their lifetime resources in choosing their current consumption and saving, one alternative model of behavior posits that individuals have short planning horizons and use a "rule of thumb" for choosing consumption and saving. A second alternative is that consumers have planning horizons that extend beyond their own lives and choose their consumption and saving to allow for bequests for their children. Yet a third alternative is that consumers are forward-looking as the lifecycle model assumes, but that they save just enough to reach some "target" level of wealth that is independent of the interest rate. Both introspection and evidence suggest that each of these models applies to some share of the population, and the response of those individuals to interest-rate changes is generally different from the response of lifecycle consumers.[1] The uncertainty surrounding these issues compounds the ambiguity of the lifecycle model, and makes it clear that theory alone cannot provide a value for the interest elasticity of saving.

Economists have tried to estimate the elasticity using data on total personal saving and interest rates, but these estimates suffer from a number of serious problems. In the end, the estimates are simply too sensitive to small changes in estimation technique to be very useful. As a consequence, the most compelling approach for determining the interest elasticity of saving is probably an indirect one.[2] This indirect approach combines models of people's behavior with estimates of certain features of their preferences and economic environment. The estimates of these features are more reliable than estimates of the interest elasticity of saving (for reasons explained below), so if the models capture the crucial elements of people's actual decision-making, this indirect approach can generate useful information about the interest elasticity of saving.

This paper pursues this indirect approach, although with two important limitations. First, the paper considers only the saving decisions of households, although direct household saving represented less than half of total private saving in the U.S. during the 1980s. The remaining

[1] Campbell and Mankiw (1989), Carroll and Summers (1991), Poterba (1994), and Bernheim (1996) all emphasize the importance of considering a variety of models of saving behavior.

[2] For a similar view, see Wilcox (1993).

part of private saving was retained earnings by businesses and contributions to defined-benefit pension plans by businesses and governments.[3] Under certain conditions, this other half of private saving can be described as simply one form of household saving, and the analysis in this paper would suffice as an analysis of the interest elasticity of all private saving. More generally, however, a separate analysis of the interest elasticity of business and pension saving is needed. A second important limitation of the paper is its focus on the short-run interest elasticity of saving rather than the long-run elasticity. I explain in the second section that the long-run effects of interest-rate changes are considerably more complicated than the short-run effects, so the paper leaves the long-run issues largely to the side.

The paper considers several alternative models of behavior. In particular, it examines the effect of interest-rate changes on the consumption and saving of people who follow the lifecycle model, people who plan to leave bequests, people who save to reach a fixed target, and people with very short planning horizons.

The analysis yields two main conclusions. First, economists' understanding of the response of saving to changes in interest rates is quite limited, despite a large volume of research on the topic. Different models of consumer behavior imply different magnitudes for the interest elasticity of saving, and even different signs. Each model probably describes the behavior of some people, and it is not clear which model best characterizes the behavior of the "average saver." Thus, it is simply not possible to provide a precise estimate of the interest elasticity of saving with any confidence. Despite the uncertainty, however, the models that likely describe the behavior of the people who account for most of aggregate saving imply positive interest elasticities. Thus, the paper's second conclusion is that the short-run interest elasticity of saving is probably positive.

The paper is organized as follows. The second section explains the relationship between interest rates and consumption and saving that is predicted by the lifecycle model. The third

[3] See Bosworth, Burtless and Sabelhaus (1991).

section reviews evidence about two aspects of individuals' preferences that affect that relationship, and the fourth section uses this evidence to quantify the interest elasticity of saving in the lifecycle model. The fifth, sixth, and seventh sections turn away from the lifecycle model and consider the effect of interest-rate changes on the consumption and saving of people who plan to leave bequests, who save to reach a fixed target, and who have very short planning horizons, respectively. The eighth section reviews other evidence on the interest elasticity of saving, and the ninth section briefly concludes.

2. The Effect of Interest-Rate Changes on Lifecycle Consumers: The Theory

The lifecycle model assumes that people are forward-looking: when choosing their current consumption and saving, they consider not only their current income and desires, but their future income and desires as well. Specifically, the model posits that a person chooses consumption at each point in his or her life to maximize lifetime well-being (utility) subject to the constraint imposed by lifetime income (the intertemporal budget constraint).

In this model of behavior, the effect of interest-rate changes on consumption and saving can be decomposed into three pieces. I present the basic logic of these pieces first, focusing on the effects in the short run. Then I discuss why the long-run effects are more difficult to analyze, which is the primary reason that this paper devotes less attention to them. In both cases, I consider only unexpected, permanent changes in the interest rate.[4] Finally, I evaluate the lifecycle model used here, briefly mentioning some complicating factors that are excluded.

Numerical formulas for the strength of the three effects are derived in Appendix A. These formulas confirm the central message of this section, which is that an increase in interest rates

---

[4] The interest rate generally changes as a result of other changes in the economy. These other changes may have separate effects on consumption and saving, but this paper studies only the direct effect of the interest-rate change.

may either raise or lower the saving of a lifecycle consumer depending on his or her preferences and economic environment. In the fourth section of the paper I show how this theoretical ambiguity is resolved under different assumptions about preferences and the environment.

Basic Logic

Two aspects of preferences are important for understanding the effect of interest-rate changes on lifecyclers' consumption and saving. One aspect is a person's *intertemporal elasticity of substitution* (or elasticity of intertemporal substitution), which reflects his or her willingness to substitute consumption between different points in life. The smaller is the intertemporal elasticity of substitution, the less a person is willing to make such a substitution. The intertemporal elasticity of substitution helps to determine the interest elasticity of saving but is not identical to it. The other aspect is a person's *rate of time preference*, which represents his or her preference for present versus future consumption. Someone with a 5 percent rate of time preference will give up $1 of consumption today only if he or she can get at least $1.05 tomorrow in exchange. The rate of time preference is essentially a measure of patience. Both of these aspects of behavior may vary across individuals, but population averages can be estimated, as discussed in the third section of the paper.

<u>Three effects</u>. For concreteness in the following discussion, I consider the effect of a permanent, unanticipated *increase* in the interest rate. The first effect of an interest-rate increase is to increase the amount of future consumption that is gained by forgoing a dollar of consumption today. By making today's consumption more costly relative to tomorrow's consumption, the interest-rate increase encourages people to consume less today and save more. This effect is called the *substitution effect*, because it involves substituting today's consumption for tomorrow's, assuming a fixed level of overall economic well-being.

The strength of the substitution effect depends on an individual's intertemporal elasticity of substitution and rate of time preference, the number of years that the individual has left to live, and the interest rate. The crucial parameter determining the strength of the effect is the intertemporal elasticity of substitution: a larger elasticity implies a larger substitution effect and an elasticity of zero implies no substitution effect at all.

The second effect of an increase in the interest rate is to lower the present discounted value (PDV) of people's planned future consumption. In other words, higher interest rates imply that fewer current dollars are needed to fund a given amount of future consumption. Planned future consumption is thus less expensive, making people better off in a lifetime sense, and leading them to consume more today and save less. This effect is called the *income effect*, and it works in the opposite direction of the substitution effect.[5]

The third effect of an increase in the interest rate is to lower the present discounted value of people's expected future income. To start, there will be a decline in the PDV of future labor and defined-benefit pension earnings, which is people's *human wealth*. There will also be a decline in the PDV of future capital income for individuals who have accumulated certain kinds of assets; this represents a downward revaluation of *financial wealth*. The revaluation of financial wealth is somewhat complicated, and will be considered more carefully below. Both of these changes make people worse off in a lifetime sense and lead them to consume less today and save more. The combined effect is called the *wealth effect*.[6]

The wealth effect does not appear in a stylized version of the lifecycle model that assumes that individuals are born with no financial wealth and that all income is earned in the first period of life so that there are no future earnings to be revalued. But the wealth effect can be quantitatively important in a more realistic model where income is earned in several periods, as stressed by Summers (1981) and demonstrated below.

---

[5] Wilcox (1993) refers to this effect as the "PDV effect," noting that it is not exactly the same as the Fisher or Slutsky income effects. Wilcox's terminology could be confusing, however, because the wealth effect discussed next is based on a PDV calculation as well. I use the more traditional term "income effect."

[6] Deaton (1992) terms this effect the "human capital effect", but that ignores the possible revaluation of financial wealth. I follow the terminology of Wilcox (1993) here.

Figure 1 presents Fisher diagrams that illustrate the effect of an interest-rate increase on an individual's budget constraint. The diagrams decompose the overall movement of the budget constraint into the substitution, income, and human wealth effects.

More about the revaluation of financial wealth. An individual's financial wealth equals the difference between his or her assets and liabilities. In considering the revaluation of financial wealth, one needs to distinguish between assets and liabilities whose value changes when interest rates change, and assets and liabilities whose value does not change.[7] A simple example in the former category is a bond. The future income flows from a bond--the coupon payments and repayment of principal--are not affected by a change in interest rates. As a result, the present value of those flows--the price of the bond--*is* affected.[8] A simple example in the latter category is a bank account. The present value of a bank account--the balance in the account--is unaffected by a change in interest rates, while the future income flows are affected.

Individuals hold financial wealth that falls into both categories. For example, the prices of stocks and bonds fall when interest rates rise, and the value of debt with a fixed interest rate--such as fixed-rate mortgages and most auto loans--falls when interest rates rise. At the same time, the value of bank accounts and debt with a floating interest rate--such as adjustable-rate mortgages and credit card balances--does not change when interest rates change. Thus, the effect of interest-rate changes on the wealth of particular individuals depends on the types of financial assets and liabilities that they hold.

For the country as a whole, however, financial wealth consists primarily of assets that are revalued when interest rates change. A bank account is an asset to its owner but a liability to someone else; a fixed-rate mortgage is a liability to the homeowner but an asset to someone else.

---

[7] Wilcox (1993) calls these two types of wealth "perpetuity-type assets" and "fixed-price assets."

[8] Note that this description applies equally well to labor income flows and the value of human wealth.

After netting out these items, aggregate wealth consists of physical capital, claims against foreigners, and perhaps claims against the government.[9] The ownership of physical capital or government bonds represents a claim to a stream of future payments that is largely unaffected by a change in interest rates; instead, the present value of those claims is affected. Thus, the simplest way to calculate an aggregate response of saving to interest-rate changes is to assume that all wealth is revalued. This method is inappropriate, however, if the types of wealth that people hold is related systematically to their propensity to consume out of wealth revaluations. I return to this issue in the fourth section of the paper when I present some aggregate calculations of the interest elasticity of saving.

Note that the financial wealth effect is captured by traditional aggregate consumption functions, in which consumption depends on income and wealth. In fact, this is the primary channel through which interest rates affect consumption and saving in many large-scale macroeconometric models. Mauskopf (1990) reports that in the MPS model used by the Federal Reserve Board, 28 percent of the effect on total spending of changes in the federal funds rate occurs through the effect of wealth revaluations on consumer spending.[10]

This entire discussion of the financial wealth effect examines the *short-run* effect on saving of a change in the interest rate. Over time, the effect of revaluing financial wealth disappears as current wealthholders die. Unborn generations do not suffer a negative wealth effect from an increase in the interest rate because they will pay lower prices for capital goods which will offset the lower PDV of the flow of income from these goods. I return to the issue

---

[9] Claims against the government are actually claims against future taxpayers, meaning ourselves and our descendants. The hypothesis that people therefore behave as if government bonds were not "net wealth" is known as "Ricardian equivalence." There is little consensus among economists as to whether this hypothesis is a useful approximation to reality; see Barro (1989) and Bernheim (1989) for opposing views.

[10] These large-scale models may not use an appropriate marginal propensity to consume (MPC) out of wealth, however. An MPC estimated using historical variation in aggregate consumption and wealth will be the average MPC for past shocks to wealth. The same MPC will not necessarily apply to shocks to wealth caused by a change in the interest rate.

of long-run effects shortly.

Combining the effects. The overall change in an individual's consumption in response to a change in the interest rate can be determined by combining the substitution, income, and wealth effects. Appendix A demonstrates that the substitution and income effects depend on the parameters of the utility function, as well as on the initial level of the interest rate and the number of years left to live. In contrast, the wealth effect depends on an individual's preferences as well as a number of features of the economic environment.[11]

I assume that an individual's income is unaffected by a change in the interest rate, and therefore that the change in her saving equals the negative of the change in her consumption. It is true that an increase in the interest rate will raise the income of individuals who hold bank accounts, but it will lower the income of bank owners by an equal amount. Aggregate income must be unchanged because the productive capacity of the economy is unchanged (in the short run). And the redistribution of income has no additional effect on the behavior of lifecycle consumers after allowing for the revaluation of wealth that has already been discussed.

The change in personal saving for the entire economy is simply the sum of the saving responses of all of the nation's households. Because households differ in the economic environment they face, their willingness to substitute consumption over time, their preference for consumption today relative to consumption in the future, and even the model that underlies their decisions, aggregation is an interesting and challenging problem. A later section of the paper

---

[11] At the same time, the income and wealth effects have quite similar logical structures because both arise from changes in the discounting of future dollar amounts in the intertemporal budget constraint. A given change in the interest rate will produce a bigger change in the present value of dollar amounts that are further in the future. Thus, if an individual earns her income before she does her consumption, the present value of consumption is more affected by interest-rate changes than the present value of income is, and the income effect will dominate the wealth effect. (Note that the income stream represented by a bond does not end at the time that an individual retires or dies. Thus, the income of older bondholders may effectively occur *after* their consumption, even though their wealth is positive. In this case, the wealth effect can be stronger than the income effect.)

presents some calculations of this sort, although the difficulty of the problem implies that the results should be viewed with great caution.

The Short Run vs. the Long Run

The preceding analysis applies to the effect of interest-rate changes on consumption and saving in the short run. The long-run effect of interest-rate changes is more complicated, as the following paragraphs explain, so this paper focuses on the short-run interest elasticity of saving.

For studying the interest elasticity of saving, the crucial distinction between the short run and the long run is whether national income can be assumed to be fixed. In the short run, a country owns an essentially fixed quantity of the factors of production, and if we abstract from business cycles, the country's income is essentially fixed as well. In the long run, however, the quantity of these factors--and thus income--depends on the country's saving over time. In fact, it is the change in the stock of wealth that is most interesting in the long run, not the change in the flow of saving.

To see why the long-run change in income complicates matters, consider two possible causes of an increase in the interest rate. Suppose first that the U.S. is a small open economy, and the world interest rate rises. For every extra dollar of private saving in a year, private wealth would increase by one dollar, and since there is no change in the government budget, national wealth would rise by a dollar as well. Private income and national income would both increase by one dollar times the marginal product of capital (the interest rate). In this case it is straightforward to calculate the long-run change in private wealth, as the formula for wealth-holding in Appendix A automatically includes the effect of higher income over time.

Now suppose instead that the U.S. is a closed economy, and a rise in the government's budget deficit increases interest rates. Every extra dollar of private saving would still raise private wealth by one dollar, but it would not raise national wealth by a dollar because of the increased public debt. In general, private wealth would rise over time, but national wealth would fall, and so would national income. Private income would increase because of the increase in private wealth, but tax payments would increase as well (to finance the increased debt service); the net effect would be a decline in after-tax private income that matches the decline in national income.

In this case the formula in Appendix A would calculate the long-run change in private wealth incorrectly unless there was an explicit adjustment for the decline in after-tax private income. The amount of this adjustment depends on the size of the debt increase that produced the rise in interest rates, but that depends in turn on the elasticity of demand for capital as well as the elasticity of supply. In other words, a given increase in the interest rate corresponds to different increases in government debt, depending on the elasticity of private demand for capital. The elasticity of capital demand can be derived from the aggregate production function, but this requires a general equilibrium model of the economy and lies well beyond the scope of this paper.

Evaluation

The third section of the paper reviews the empirical evidence regarding the key parameters of people's utility functions, and the fourth section uses these estimates in the framework presented above to examine the response of consumption and saving to changes in interest rates. Before going on, however, it is worth considering the value of the lifecycle model presented here for understanding people's behavior. A full evaluation of the model is clearly beyond the scope of this paper, but two issues need to be discussed briefly. First, does the central insight of the lifecycle model--that people make current decisions based on their lifetime circumstances--appear to be valid? And second, does the simplified lifecycle model used in this paper capture the decisive aspects of those lifetime circumstances and decisions?

On the first question, one piece of evidence that the lifecycle model does not characterize the behavior of many people is that many households hold essentially no financial assets. For example, over one-third of households had less than $10,000 in financial assets in 1992, according to the 1992 Survey of Consumer Finances. Thus, a significant share of the population does almost no saving, aside from (perhaps) mortgage payments and pension contributions. But this does not imply that the lifecycle model is inappropriate for studying aggregate saving. The saving decisions of people who do little or no saving presumably have little impact on aggregate saving, so it will not matter very much for aggregate purposes just how they make those saving decisions. If the people who are doing significant amounts of saving follow the lifecycle model, then we can use the model to learn about the determinants of aggregate saving and about the role of the interest rate in particular.[12] For completeness, however, the seventh section of the paper examines the effect of interest-rate changes on the saving of individuals who are not as forward-looking as lifecycle consumers.[13]

There is also evidence that some people are *more* forward-looking than implied by the lifecycle model, in that their saving decisions reflect concern for their children as well as for themselves. These people save not only to finance their own future consumption but also to finance bequests to their children. The fifth section of the paper explores the effect of interest-rate changes on the saving of these bequest leavers. Finally, there are other people who are forward-looking as the lifecycle model assumes, but save just enough to achieve a given level of consumption later. The sixth section of the paper explores how these so-called "target savers" respond to changes in interest rates.

---

[12] Deaton (1992) argues that if consumption is always close to income for many people, then "the prima facie supposition must be that interest rates are not very important" (p. 75) in determining consumption and saving. But consumption is not that close to income for people who *are* doing the saving in the economy, so there is no reason to presume that the *savers'* decisions are unaffected by interest rates or other factors.

[13] Some people with very little wealth may be trying to follow the lifecycle model, but are prevented from doing so by an inability to borrow against their expected future income. Rising labor income paths and significant annuity streams imply that many people's labor income is tilted toward the end of their lives; people in this situation who are fairly impatient may accumulate no financial wealth until fairly late in their working lives. Revaluations of human wealth can only affect the current consumption of these people if they borrow against their human wealth, but the credit markets make this type of borrowing almost impossible.

The second question raised above is whether the specific version of the lifecycle model presented here omits important features of the lifecycle decision process. Two omissions deserve mention. First, the analysis assumes that people live in a world of certainty regarding both the length of their lives and their future incomes.[14] In more realistic models that allow for income uncertainty, people have an additional, precautionary, motive for saving: they save to hedge against bad realizations of future income. In models that allow for lifespan uncertainty, people save to protect against living an unexpectedly long time. These precautionary motives create a number of important differences in people's behavior, including the interest elasticity of their saving. Engen (1992) shows that the interest elasticity in an uncertain world is substantially smaller than in a corresponding world of certainty, although the magnitude of the reduction is sensitive to auxiliary assumptions regarding the rate of time preference and the amount of uncertainty.[15]

Unfortunately, the empirical significance of precautionary saving remains unclear. On one hand, allowing for a significant amount of precautionary saving can help to explain some important patterns in consumption and saving at both the household and aggregate levels.[16] On the other hand, many of these patterns have alternative explanations, and direct evidence on the importance of precautionary saving is mixed.[17] On balance, lifecycle models with certainty clearly overstate the true interest elasticity of saving, but the amount of that overstatement is

---

[14] Interest rates change once, unexpectedly, but both before and after the change, people behave as if they face no uncertainty about the future.

[15] The introduction of uncertainty reduces the interest elasticity of saving by an average of 70 percent in the examples that Engen presents. At the same time, precautionary saving represents about 80 percent of total saving in those examples, which seems unrealistically high.

[16] Engen (1992) and Hubbard, Skinner and Zeldes (1995) discuss precautionary saving and household behavior; Caballero (1990) does the same for aggregate behavior.

[17] For example, the wealth-income ratio in the U.S. is higher than is predicted by lifecycle models with standard parameter choices. This phenomenon can be explained by the existence of precautionary saving, but it can also be explained by altruistic bequests. Carroll (1992) and Dynan (1993) present opposing direct evidence on the strength of the precautionary saving motive.

unknown. This is clearly an important area for future research.

A second omission from the lifecycle model presented here is durable goods. An increase in the interest rate raises the user cost of durable goods, encouraging people to consume less durable goods and more nondurable goods and other services. If investment in durable goods is viewed as saving, then this effect is simply part of the overall changes in consumption and saving analyzed in this paper. If investment in durable goods is not viewed as saving, however, then the shift away from durable goods just described represents a change only in the composition of consumption.[18]

To summarize, the lifecycle model clearly does not describe everyone's consumption and saving decisions, although it may describe the source of most aggregate saving. Further, this paper considers the interest elasticity of saving under several alternative models of individual behavior. At the same time, the specific lifecycle model used here ignores the uncertainty facing households in the real world; this simplification greatly improves the tractability of the analysis, but it produces an overstatement of the interest elasticity of lifecyclers' saving.

3. Empirical Evidence on Individuals' Preferences

The previous section explains that two characteristics of an individual's preferences affect the response of his or her saving to changes in the interest rate. The more important characteristic is the intertemporal elasticity of substitution, and most of this section is spent reviewing empirical estimates of this parameter. The less important characteristic is the rate of time preference; the more limited evidence on this parameter is also discussed in this section. To preview the conclusions, the intertemporal elasticity of substitution appears to lie between zero and one-half, while the rate of time preference is probably positive but could actually be negative.

[18] See Mankiw (1987) and Wilcox (1993) for analysis of this issue.

The Intertemporal Elasticity of Substitution

Many studies have produced estimates of the intertemporal elasticity of substitution. Most of the recent estimates are based on data on consumer spending, analyzed using the Euler equation framework pioneered by Hall (1978).[19] The Euler equation for the lifecycle model of consumption, which shows how an individual's optimal consumption evolves over time, is:

$$\frac{E(\Delta C_t)}{C_t} = \mu + \sigma \, (E(r) - \delta) \ ,$$

where C_t is consumption in period t, σ is the intertemporal elasticity of substitution, δ is the rate of time preference, r is the interest rate, and $E()$ denotes the expected value. This equation implies that a 1 percentage point decline in the expected interest rate, say from 4 percent to 3 percent, will reduce the rate of growth of consumption by σ percentage points. Therefore, one can estimate σ by regressing the percent change in consumption on the expected real return on saving, possibly including other explanatory variables as well.[20]

The following review of the empirical literature on the intertemporal elasticity of substitution discusses the two most widely-cited papers based on aggregate data and a number of papers based on household-level data. Table 1 summarizes these papers' estimates of this parameter.

Estimates derived from aggregate data. Hall (1988) examines the behavior of aggregate consumer spending, using different time periods and different measures of the expected rate of return. He finds no evidence of a significant positive elasticity of substitution. Hall notes that his results are not surprising given that the U.S. has experienced large changes in the interest rate

[19] See Randolph and Rogers (1995) for a discussion of other data that have been used to estimate this parameter.

[20] In fact, most studies that estimate σ focus principally on the role of these additional variables.

over time but only small changes in the growth rate of aggregate consumption. Together, these facts suggest that consumption is not very responsive to changes in the interest rate and thus that the intertemporal elasticity of substitution is small.[21]

Campbell and Mankiw (1989) also examine aggregate consumption data, extending Hall's framework to allow for some households who choose consumption based on a "rule of thumb" rather than the lifecycle model. They estimate that these households receive almost half of total income, suggesting that their presence significantly alters the expected dynamics of aggregate consumption. For the traditional lifecycle consumers, Campbell and Mankiw estimate elasticities of substitution between 0 and 0.2.

Estimates derived from household data. One of the earliest papers to use cross-sectional variation in interest rates to estimate the intertemporal elasticity of substitution is Shapiro (1984). His estimates are very large, but they have enormous standard errors, so they are simply not informative.

Zeldes (1989) estimates the intertemporal elasticity of substitution as part of his exhaustive study of liquidity constraints. He divides the Panel Study of Income Dynamics (PSID) sample into households with a high wealth-income ratio (who are unlikely to be liquidity-constrained) and households with a low wealth-income ratio (who may be liquidity-constrained). For different measures of wealth, Zeldes finds consistent evidence of liquidity constraints but widely different estimates of the elasticity of substitution, ranging from -1.46 to 1.92.

Runkle (1991) studies liquidity constraints with the PSID as well, but in contrast with Zeldes, Runkle finds no evidence of such constraints. It is generally difficult to reconcile Runkle's findings with Zeldes's; Deaton (1992) suggests that the difference stems from Runkle's

---

[21] Lucas (1990) reaches the opposite conclusion by applying similar logic to cross-country data. Lucas states that the growth rate of consumption varies widely *across countries* but the rate of return does not, suggesting a large response of consumption to interest rates and a large elasticity of substitution.

exclusion of more than 80 percent of the observations in Zeldes's sample. Runkle estimates a fairly large elasticity of substitution--roughly 0.6--for households with a high wealth-income ratio, and a smaller elasticity--roughly 0.3--for households with a low wealth-income ratio. In both cases, the standard errors are fairly small.

Lawrance (1991) investigates whether the rate of time preference differs systematically between rich and poor households. She divides the PSID sample by income and finds large elasticities of substitution--0.81 and 1.80--with relatively small standard errors for households with income above the median. Households with below-median income appear to have a smaller and less statistically significant elasticity. The difference in estimated elasticities between rich and poor households is potentially important, as the rich do a large fraction of total saving, and thus largely determine movements in the aggregate saving rate.

Dynan (1993) uses data from the Consumer Expenditure Survey (CEX) to determine the strength of the precautionary saving motive. She estimates a quite small elasticity of substitution, around 0.1, but with a large standard error.[22]

Attanasio and Weber (1993; hereafter AW) investigate the effects of aggregation on consumption Euler equations, with particular emphasis on the intertemporal elasticity of substitution. Using British National Accounts data, they estimate an elasticity around 1/3, which is somewhat larger and has a smaller standard error than the Campbell-Mankiw estimate for the U.S. Then AW turn to the British Family Expenditure Survey (FES), which provides a long time series of household-level data. They construct separate time series for the average consumption of all households and of households in three different age cohorts. For all households, the estimated intertemporal elasticity of substitution is around 1/3, as in the aggregate British data. For the young and middle-aged cohorts, the estimated elasticity is roughly twice as large, and for

[22] As Dynan points out, however, her small estimates of the coefficient of relative prudence imply (for a broad class of utility functions) a small coefficient of relative risk aversion and a large elasticity of substitution.

the older cohort, it is fairly small (although not statistically different from the other estimates).

In another paper, Attanasio and Weber (1995) argue that previous rejections of the lifecycle model can be traced to simplifying assumptions used in other studies. They emphasize the importance of examining household-level data on the consumption of all commodities, by illustrating the biases that can arise in estimates using either aggregate data or household data on food consumption only. They also show the value of allowing a household's consumption to change over time in response to changes in household composition and labor supply. Attanasio and Weber's preferred specifications (using a synthetic panel constructed from the CEX) produce estimates of the intertemporal elasticity of substitution between one-half and two-thirds.

Finally, Attanasio and Browning (1995) use data from the FES to study the sensitivity of consumption to income over both the lifecycle and the business cycle. Their parameterization of preferences allows both the rate of time preference and the intertemporal elasticity of substitution to depend on household characteristics. Attanasio and Browning find that the elasticity of substitution increases with consumption, which is consistent with Lawrance's result. The authors do not report an elasticity estimate, however, which would be a nonlinear function of several parameters and household variables.

Evaluation. Unfortunately, there are good reasons to be skeptical of estimates of the intertemporal elasticity of substitution based on either aggregate or household data. Consider the aggregate approach first. The Euler equation is derived from the optimization problem of a *single* individual or household, and the relationship it reveals between individual consumption growth and the interest rate does not necessarily apply at the aggregate level. For example, Deaton (1992) explores the assumptions necessary for aggregation and concludes that "it is quite unsafe to make any inference about intertemporal substitution from representative agent models" (p. 68).

Studies based on household-level consumption data avoid the aggregation problem, but introduce other problems of their own. First, the data available for U.S. households cover only a small number of time periods (five quarters in the Consumer Expenditure Survey) or a small part of consumption (food consumption in the Panel Study of Income Dynamics). These limitations reduce the accuracy of the estimated Euler equations.[23] Second, the studies generally rely on variation in the rate of return across households that is caused by differences in marginal tax rates. But tax rates are correlated with income and thus with households' tastes and other characteristics. If household saving depends directly on any of these other characteristics, then this heterogeneity could bias estimates of the elasticity of substitution.[24] Third, the studies apply different specifications to separate data sets using assorted sample selection procedures, which makes it hard to reconcile all of their findings. As a result, Deaton concedes that microeconomic studies of consumption "have not produced the sort of widely agreed stylized facts" that the macroeconomic literature has (p. 136).

With these caveats in mind, it is clear that the existing empirical evidence does not allow us to identify the intertemporal elasticity of substitution with any precision. Economic theory says that the elasticity must be greater than zero.[25] The empirical evidence implies that the elasticity is less than one and probably less than one-half. More specifically, a crude average of the estimates reported in Table 1 (based on equal weights for each estimate within a study, and then equal weights for each study aside from Shapiro's) is 0.37. I use an estimate of 0.33 in the calculations in the next section of the paper.

A value of 0.33 for the intertemporal elasticity of substitution is consistent with the values

---

[23] Also, these data sets do not contain many observations of households at the very top of the income distribution. If the intertemporal elasticity of substitution is quite skewed in the population (which would be consistent with the skewness of saving), then the estimates discussed here may miss the larger values.

[24] Income may also be correlated with pre-tax rates of return. Feldstein (1994) argues that high-income households may earn a higher *pre-tax* return on their investments than other households, and thus may have a higher after-tax return despite their higher tax rates.

[25] A negative value for the elasticity would imply that the marginal utility of income does not decline with income (utility is not concave), and a value of zero would imply that people are infinitely risk averse (for a broad class of utility functions).

selected by many previous researchers for simulation exercises and analytical calculations. Auerbach and Kotlikoff (1987) use 0.25 in their book on dynamic fiscal policy, Lucas (1990) uses 0.5, King and Rebelo (1990) use 1.0, and Gravelle (1991) uses 0.25. Deaton (1992) chooses values of 0.33 and 0.5, Jones, Manuelli and Rossi (1993) choose 0.67, and Hubbard, Skinner, and Zeldes (1995) select 0.33 as their base case.

The Rate of Time Preference

There is little direct empirical evidence regarding the rate of time preference. The Euler equation methodology that enables researchers to use data on consumer spending to estimate the intertemporal elasticity of substitution does not provide such a straightforward procedure for estimating the rate of time preference.

The difficulty with the Euler-equation approach can be seen by rewriting the Euler equation from above as follows:

$$\frac{E(\Delta C_t)}{C_t} = \sigma \; E(r) + (\mu - \sigma \; \delta) \; .$$

The rate of return varies in an observable way across time and across individuals, so estimates of σ can be based on the correlation between that variation and the variation in the consumption growth rate. The second term on the right side of the equation is the combination of several unobservable parameters, so it is estimated as a constant. If μ equalled zero, one could back out an estimate of δ using the estimate of σ from the first term; in fact, μ is not zero but depends on the variance of people's errors in forecasting consumption growth. This variance can be estimated, but the estimates are quite imprecise. As a result, most of the papers that use Euler equations to estimate the rate of time preference have focused on differences in the rate across groups rather than the average level in the population.[26]

[26] See Lawrance (1991) and Dynan (1994).

Economic theory also has little to say about the rate of time preference. The basic axioms of consumer choice do not specify whether people should discount the future relative to the present (implying a positive rate of time preference), or weigh the future more heavily (implying a negative rate of time preference). One argument in favor of a positive rate of time preference is that negative rates are incompatible with infinite-horizon models: infinitely-lived individuals who are patient would defer consumption so much that the economy would not achieve a steady-state capital stock or balanced growth path.[27] In an economy with many generations, however, negative rates of time preference *within a person's lifetime* do not pose a problem even if the generations are altruistically linked.[28]

Because standard theory and empirical techniques provide little information about the rate of time preference, the available evidence regarding this parameter is fairly informal. One indication that most people are fairly impatient (that is, have a substantially positive rate of time preference) is that most people do little or no saving. On the other hand, most people can expect rising earnings over their lifetime, and the calculations in the next section show that even patient people in that situation may save little until well into their working lives. Further, the behavior of people who do little saving may simply not be very important for the issue addressed in this paper, as argued at the end of the second section.

Another indication of the average rate of time preference may be gleaned by examining total consumer spending. Suppose that the Euler equation applies to spending per capita, and that the unobservable factors that make it difficult to isolate δ in household data can be ignored at the aggregate level. Real per-capita consumer spending grew by roughly 2 percent per year between 1950 and 1994; if σ equals 0.33 and the real rate of return has averaged 4 percent, then the implied rate of time preference is -2 percent.

---

[27] See Olson and Bailey (1981).

[28] Negative rates of time preference *across generations* would represent an extremely strong form of intergenerational altruism, and could create the same instability as in the infinite-horizon world.

A final technique for acquiring information about the rate of time preference is to ask people (the right) questions directly. Loewenstein and Prelec (1991) argue that impatience and the desire for improving conditions "are simultaneously present within a single individual" (p. 348), and that the context of a decision determines the dominant motive. They report on a number of surveys they conducted, and conclude that isolated intertemporal choices encourage impatience, but choices that are part of a "coherent sequence" (p. 351) lead to patient behavior.

Barsky, Juster, Kimball and Shapiro (1995) explain that people's rates of time preference and intertemporal elasticities of substitution can be deduced from these individuals' choices among alternative hypothetical time paths for consumption. Barsky et al construct a set of alternative paths that were presented to a small subset of the participants in the Health and Retirement Survey. They find that more of the people surveyed displayed a negative rate of time preference than a positive rate; in fact, the average rate of time preference is roughly -4 percent.[29] Unfortunately, there are only 198 responses to these questions, of which 30 percent are inconsistent with utility maximization in some way and are dropped from the analysis. The remaining respondents are not only few in number but may be confused by the questions as well.

In conclusion, there is little firm evidence concerning the sign of the rate of time preference, although most economists find the concept of impatience sufficiently plausible to assume a positive rate.[30] The calculations in the next section of the paper use rates equal to -.02, 0, and 0.02, although more weight should be placed on the positive end of that range.

4. The Effect of Interest-Rate Changes on Lifecycle Consumers: Illustrative Calculations

[29] The average intertemporal elasticity of substitution among these respondents is 0.18. The implied rate of time preference equals the negative of desired consumption growth at a zero interest rate (0.78) divided by 0.18. See Barsky et al, Table 7.

[30] For example, see Blanchard and Fischer (1989, p. 39).

The second section of the paper describes the way in which a lifecycle consumer would adjust his or her consumption and saving in response to a change in the interest rate. This section uses the formulas derived in Appendix A to quantify these adjustments for some illustrative consumers. The illustrations use alternative assumptions about people's preferences--based on the evidence reviewed in the previous section--and the economic environment that people face. It will be evident that the response of saving to interest-rate changes varies significantly across the examples. The section then considers how individual examples like these can be used to predict the interest elasticity of aggregate saving. As might be expected, such aggregation is a complex task, and I address only some of the relevant issues. The section concludes with a brief discussion of previous related work.

It is important to emphasize that the calculations in this section are intended to be illustrative, not definitive. To start, many people may not follow the lifecycle model in choosing their consumption and saving; the interest-rate responses of these people are explored in later sections of the paper. Further, the calculations ignore people's uncertainty about their lifespan and future income, which reduces the interest elasticity of saving in a lifecycle model. And even within the certainty case, the calculations explore only a limited range of preferences and economic environments. In summary, the following calculations offer suggestive evidence regarding the interest elasticity of saving and the factors that influence it, but they do not produce an authoritative estimate of the elasticity.

A Base Case

The calculations throughout this section of the paper are based on the following assumptions. People begin working and consuming at age 25, retire at age 65, and die at age 80. Labor earnings are constant over a person's lifetime in the initial examples, but in later examples they increase with time. After retirement, people receive 30 percent of their pre-retirement labor earnings through Social Security and private pension plans.[31] People can hold their accumulated wealth in two forms: "bank accounts," which are not revalued when the interest rate changes, and shares of "physical capital," which are revalued. The stream of payments from a unit of physical capital declines over time as the capital depreciates; I assume a depreciation rate of 7 percent.[32]

The initial interest rate is assumed to be 4 percent; because there is no explicit treatment of inflation here, the rate should be interpreted as the real return on financial wealth. The calculations examine the immediate effect on saving of a small permanent change in this rate of return. (See the second section of the paper for a discussion of immediate versus long-run saving responses.)

Table 2 presents the results of the calculations for a convenient base case. The intertemporal elasticity of substitution, the rate of time preference, and the growth rate of wages are all assumed to be zero. The table presents the effect of the interest-rate increase on the saving behavior of a 30-year old, a 50-year old, and a 70-year old. Each individual's overall interest elasticity of saving is decomposed into the elasticity arising from the substitution effect, the income effect, and the wealth revaluation effect.

The first line of the table shows the substitution effect, which is exactly zero here because the intertemporal elasticity of substitution is assumed to be zero. The second line gives the income effect, and the third line gives the effect of revaluing human wealth. The fourth line shows the effect of revaluing financial wealth if it is held entirely in the form of bank accounts, while the fifth line shows the effect of revaluing financial wealth if it is held entirely in the form

---

[31] This 30-percent figure is too low for the population as a whole, but may be appropriate for higher-income people who do most of total saving.

[32] Dividend streams from common stock may increase over time, but those increases probably reflect the use of retained earnings to purchase new capital, not an increase in the productivity of capital already owned. According to the Flow of Funds Accounts, .07 is the average ratio from 1985 to 1994 for non-farm, non-financial corporations of (NIPA depreciation with the capital consumption and inventory valuation adjustments) to (the current cost of tangible assets).

of physical capital. The sixth and seventh lines display the total elasticity of saving if wealth is held entirely in bank accounts or in physical capital, respectively.

An increase in the interest rate would cause both the 30-year old and the 50-year old in Table 2 to reduce their saving, as the income effect outweighs the effect of revaluing human and financial wealth. The 50-year old experiences smaller income and human wealth effects than the 30-year old because she has less consumption and earnings remaining in her life, and it is the change in the present discounted value of those streams that gives rise to these two effects. But the 50-year old has accumulated more financial wealth than the 30-year old, so she experiences a larger effect from its revaluation. The 70-year old has yet smaller income and human wealth effects (the human wealth effect is not zero because of future pension income), but a larger financial wealth effect if he holds physical capital. A 70-year old holding only physical capital increases his saving (actually, decreases his dissaving) if the interest rate increases.

One way to understand the difference in response between a young person and an older one is the following. Young people expect to earn their human wealth before they do their consumption; thus, increases in interest rates relax their budget constraint, making them better off and reducing their saving. As people get older, however, their human wealth is smaller and their financial wealth larger. If this financial wealth is held as physical capital, the associated payments will occur largely after consumption, and increases in interest rates will make holders worse off (as the price of physical capital declines) and increase their saving.

The Effect of Alternative Assumptions

The following examples explore the effect on the interest elasticity of saving of alternative assumptions about utility parameters and the economic environment.

<u>The intertemporal elasticity of substitution</u>. The second section of the paper showed that 0.33 was roughly the midpoint of existing estimates of the intertemporal elasticity of substitution. In light of the difficulties in estimating this parameter, however, significantly smaller or larger

values cannot be ruled out.

Table 3 reproduces Table 2 but assumes that the intertemporal elasticity of substitution equals 0.33 rather than zero. The unsurprising result here is that the interest elasticity of saving is less negative or more positive for people of all ages and holding either type of wealth. In fact, for people who hold their wealth in the form of physical capital, the elasticity of saving is now positive at all ages.

Similar calculations are easy to perform for other values of the intertemporal elasticity of substitution. Calculations based on a value midway between zero and 0.33 (not shown) produce saving elasticities that are closer to the ones in Table 3 than those in Table 2.

<u>The rate of time preference</u>. The second section of the paper concluded that the rate of time preference could be either positive or negative, because economic theory and existing empirical evidence do not provide much guidance. Nevertheless, most economists consider a positive rate of time preference more likely. Tables 4 and 5 reproduce Table 3 but allow the rate of time preference to be -0.02 and 0.02, respectively.[33]

A negative rate of time preference--meaning that people are more patient than in the base case--makes the elasticity of saving slightly less positive or more negative for all of the illustrative individuals. A positive rate of time preference makes the elasticity of saving more positive or less negative for all of the individuals, and substantially so for some of them.

Despite the differences created by varying the rate of time preference, the qualitative story from Table 3 is largely unaffected. People who hold physical capital generally have a positive interest elasticity of saving, and people who hold bank accounts generally have a negative interest

[33] The intertemporal elasticity of substitution is held at 0.33. A person's rate of time preference has no effect on his interest elasticity of saving if his intertemporal elasticity of substitution is zero, because he would set his consumption equal in all periods regardless of his degree of patience.

elasticity of saving. There is one way, however, in which varying the rate of time preference can have a significant influence on the elasticity: very impatient individuals are more likely to want to borrow against their income and are thus more likely to be affected by borrowing constraints. The effect of binding borrowing constraints is considered below.

<u>Labor income growth over a lifetime</u>. Table 6 introduces earnings growth over a person's lifetime, with the calculations assuming that an individual's labor earnings increase by 1 percent per year. For the other parameters, I return to the assumptions of Table 3, with an intertemporal elasticity of substitution equal to 0.33 and a rate of time preference equal to zero.

The change in the *amount* of saving caused by the substitution and income effects is the same for individuals here and in Table 3, if their initial labor earnings are adjusted so that they have the same lifetime wealth. (This result is unsurprising because both effects are based on the stream of planned consumption, which is the same.) But the absolute value of the *elasticities* caused by these effects are larger here because initial saving is smaller. The human wealth effect is substantially stronger here, both because initial saving is smaller and because the bulk of earnings come somewhat later in life, so their present value is affected more strongly by a change in the interest rate. The financial wealth effect is weaker, however, because a person who earns more later in her life will accumulate less wealth. On balance, saving elasticities are more positive or less negative in Table 6 than in Table 3, although the difference is only significant for 30-year olds.

<u>Liquidity constraints</u>. Table 7 raises the growth rate of earnings over a lifetime to 3 percent per year, which seems more realistic for many people. This growth rate generates a steeply-sloped earnings path: earnings at age 65 are almost three times earnings at age 25. Recall also that an individual's pension is assumed to be a fixed fraction of earnings in his last year of work. This is only a slight exaggeration of the real world, where benefits from Social Security and most defined-benefit pension plans depend primarily on earnings in the last several years of work. Other parameters match those in Tables 3 and 6.

In this situation, a lifecycle consumer who does not face liquidity constraints would borrow in the first part of her life in order to maintain a level of consumption that reflects her permanent, rather than current, income. Her financial wealth would become increasingly negative for a number of years, then turn and eventually become positive, before falling again to zero at the end of her life. Her biggest indebtedness would be roughly 1.5 times her annual earnings, and her financial wealth would actually remain negative until age 52.[34]

In the real world, it is generally not possible to borrow against one's human wealth to finance a significant amount of consumption. It is sometimes possible to finance education expenses in this way using student loans, and it is possible to finance a small amount of consumption in this way using credit cards. But the person with the optimal consumption plan described in the previous paragraph would almost certainly face a liquidity constraint that would prevent her from carrying out that plan. Suppose that this person is unable to hold negative financial wealth. How would her consumption and saving respond to a change in the interest rate under this circumstance?[35]

While the liquidity constraint was binding, the person would do no saving; thus, her saving would be unaffected by a change in the rate of return.[36] Table 7 thus shows no response for 30-year olds. (The seventh section of the paper discusses people who are liquidity-

---

[34] With a rate of time preference equal to -0.02, wealth hits -0.2 times annual earnings and remains negative until age 37. With a rate equal to 0.02, wealth hits -3.4 times annual earnings and remains negative until age 58.

[35] If the person were unconstrained in her borrowing, her elasticity of saving at age 30 would be larger than shown in Table 6, because an increased human wealth effect from an interest-rate increase would offset a negative effect from the revaluation of her negative financial wealth. Her elasticity at age 50 would be smaller, however, as the increased human wealth effect was more than offset by a large reduction in the financial wealth effect.

[36] A sufficiently large increase in the interest rate could encourage this person to shift her desired consumption late enough in her life that she would no longer be liquidity-constrained. In this case she would have a positive elasticity of saving as her saving increased from zero to some positive amount.

constrained.) When this person reached age 41, however, the constraint would no longer be binding: she would begin to save, and the amount of her saving would be responsive to changes in the interest rate. At age 50, the human wealth effect is stronger than for someone with slower wage growth, but the financial wealth effect is weaker because she has been accumulating wealth for a much shorter time. By age 70, however, the effects are similar to those shown in Tables 3 and 6.

Aggregation and Heterogeneity

As the previous examples show, the effect of interest rate changes on saving is quite different for individuals in different situations, even if all are following the lifecycle model. These differences arise not only from differences in preferences, but also from differences in age, type of wealth held, the growth rate of wages, the presence of liquidity constraints, and probably other factors as well. In light of this heterogeneous response, determining the aggregate interest elasticity of saving--even conditional on this basic lifecycle model--is extremely difficult. The following paragraphs examine some of the issues that arise in aggregation, but does not try to present a definitive estimate.

Suppose first that the only difference among people were their ages. Then one could calculate an aggregate saving response for a given set of other characteristics by cumulating the responses of people of different ages. Three complications arise in this plan. First, it is not straightforward to aggregate the elasticities directly, because they apply to different levels of saving and dis-saving; it is more sensible to cumulate *changes* in saving, and then calculate the aggregate saving elasticity at the end. Second, the changes in saving for people of different ages need to be weighted to account for the growth of productivity and population across birth cohorts.[37] Younger generations in this country have higher lifetime income than older generations, and include more people, so the saving behavior of the typical young person has a

[37] Note that the growth of productivity across birth cohorts is a separate issue from the age-earnings profile of individuals in any given cohort.

larger effect on aggregate saving than the saving behavior of the typical older person. For example, if aggregate productivity grows 2 percent per year and the population grows 1 percent per year, the saving response of a 25-year old should receive a weight that is more than four times greater than the weight given to the response of a 70-year old. In the calculations reported below, I use combined productivity and population growth rates between 2 and 4 percent.

The third complication is deciding whether people hold their financial wealth in the form of bank accounts or physical capital. As explained earlier, aggregate financial wealth consists primarily of physical capital, because bank accounts are assets to some wealth-holders but liabilities to others. Thus, the simplest approach is just to ignore bank accounts (and similar assets and liabilities), and pretend that everyone holds their wealth entirely as shares of physical capital. The problem is that physical capital carries the risk of large revaluations, such as the ones that arise from interest-rate changes, and older people should be particularly averse to bearing that risk because of their higher marginal propensity to consume (MPC).[38] Therefore, older people will tend to hold their wealth in bank accounts, while young people's wealth will consist of the economy's physical capital assets less the liabilities of the older people's bank accounts. Because the old have a higher MPC, shifting the ownership of physical capital toward the young reduces the effect of wealth revaluations and thus the interest elasticity of saving. I implement this alternative approach by assuming that the share of wealth held as physical capital falls linearly as people age.

Table 8 presents some calculations of the aggregate interest elasticity of saving corresponding to the parameters used in Table 7. For a combined productivity and population growth rate of 3 percent, and declining physical capital holding over a lifetime, the elasticity is 0.65. Note that this value exceeds the elasticity for both age groups shown in Table 7, and for almost all of the age groups not shown as well. This surprising relationship arises in the following way. An increase in the interest rate encourages almost every age group to raise its saving; those increases are expressed as elasticities by dividing by the absolute value of that

[38] I am grateful to Bill Randolph for pointing this out.

group's initial saving (and by the percent change in the interest rate). If one calculated the aggregate elasticity by dividing the change in aggregate saving by the sum of *the absolute value* of each group's saving, then the aggregate elasticity would be a weighted average of the elasticities of the groups. But the true aggregate elasticity equals the change in aggregate saving divided by the sum of each group's saving, which is a much smaller number if a lot of people are doing negative saving. That is indeed the case in this example: gross saving (by young people) is 15 percent of aggregate income, and gross dis-saving (by older people) is 9 percent of income, leaving the reported saving rate of 6 percent.

Moving beyond this example, suppose next that people differ in ways other than their age. These differences could be in preferences, the age-earnings profile, the level of income, or something else.[39] How would this heterogeneity affect our ability to construct an aggregate elasticity of saving?

If the interest elasticity of saving varied linearly with respect to a certain variable, then the elasticity calculated for the average value of that variable would be correct. Unfortunately, such a linear relationship does not generally hold. One example is the effect of the intertemporal elasticity of substitution, which was discussed above. Another example is the difficulty of borrowing against one's human wealth, which implies that an economy where a few people have very high wage growth and many people have no wage growth will have fewer people pushed into a constrained optimum than an economy where all people have a small amount of wage growth. Yet a third example is that the rate of time preference may be positively correlated with income, if people who are more patient invest more in human capital. In this case, the income-weighted average rate of time preference will exceed the population-weighted average rate of

---

[39] For the utility function used in this paper, a person's interest elasticity of saving does not depend on her income: people with higher income do more saving initially *and* make a correspondingly greater change in their saving when interest rates change. Bernheim and Scholz (1993) suggest an alternative utility function for which the interest elasticity of saving is positively correlated with income. They show that if utility is a function of actual consumption less some subsistence level of consumption, the "effective" interest elasticity of saving increases with consumption.

time preference.

In general, allowing for heterogeneity is quite important for understanding saving behavior, because the distributions of saving and wealth are quite skewed.[40] Since a small fraction of the population does a large share of total saving, any response of saving to a change in interest rates would presumably be concentrated among those people. What role can data play in helping to surmount these difficulties? Data on individuals' consumption decisions has been used extensively to learn about preferences, as reviewed in the previous section. Beyond that, data on the distribution of earnings by age and other characteristics could be used to estimate more carefully the revaluation of human wealth in response to an interest rate change. And finally, household-level data on financial wealth holdings might be used to calculate the financial wealth effect more precisely.[41]

Related Research

There have been a number of attempts to calculate the interest elasticity of aggregate saving using an indirect method similar to the one used in this paper. Like the calculations presented in this section, these studies use the lifecycle model to combine estimates of behavioral parameters with auxiliary assumptions about the economic environment. Yet these studies differ from the calculations presented in this section in at least two important ways. First, the studies are usually based on models of the entire economy rather than calculations derived from a single individual's optimization problem. A natural advantage of these general equilibrium models is that they require no further aggregation. Second, the studies generally focus on changes in

---

[40] See Bosworth, Burtless and Sabelhaus (1991) and Avery and Kennickell (1991).

[41] Wilcox (1993) calculates an aggregate financial wealth effect using data on aggregate wealth from the Flow of Funds accounts compiled by the Federal Reserve Board. He first determines whether a certain category of wealth would be revalued if interest rates change; if so, he calculates the approximate amount of the revaluation and applies a commonly used MPC out of wealth to ascertain the change in aggregate consumption. One problem with this aggregate approach is that MPCs should differ across individuals.

wealth over a very long horizon rather than the immediate changes in saving presented here. In other words, the studies emphasize the steady-state change in the capital stock, meaning the change after the economy has fully adjusted to the new rate of return.

The approach followed in this paper is not as sophisticated as the approach used in these other studies, and the results presented here are clearly less informative on some important issues. At the same time, the relative simplicity of the calculations in this paper helps to highlight the roles played by different effects and behavioral parameters in the overall response of saving to interest-rate changes. Because of the difference in methods, it is not possible to compare directly the results of the previous studies with the illustrative calculations presented in this section. A few comments about some of these studies may be useful, however.

Summers (1981) analyzes the effect on the nation's capital stock of shifting away from the taxation of capital income and thereby raising the after-tax return to saving. He stresses the importance for the interest elasticity of saving of the revaluation of expected labor income. For an intertemporal elasticity of substitution (σ) equal to 0.33, a rate of time preference (δ) equal to 0.03, and an initial interest rate (r) of 0.04, Summers estimates an interest elasticity of 2.38 (Table 1, p. 536). His estimated elasticities are smaller for higher initial interest rates and for lower elasticities of substitution. Summers discusses the effect of financial wealth revaluations in evaluating empirical work on this topic, but such revaluations play no role in his estimates because the estimates are based on steady-state changes in saving. As explained earlier, the effect of financial wealth revaluations disappears eventually as current wealthholders die.

Evans (1983) argues that Summers' estimated elasticities are quite sensitive to the rate of time preference, and become much smaller for negative values of δ. For σ equal to 0.33, δ equal to 0.03, and r equal to 0.04, Evans estimates an elasticity of 2.97, quite close to Summers' estimate. But for δ equal to zero or -0.03, the elasticity falls to .56 or .41, respectively (Table 1, p. 401). Most of the decline in the elasticity is caused by a higher initial level of saving; the remainder is caused by a smaller predicted change in saving. These differences are similar qualitatively to those in Table 4, but much more dramatic quantitatively.

Choosing a lower rate of time preference plays another, indirect role for Evans in justifying a smaller estimate of the interest elasticity of saving. One way that Summers defends his preference for large values of σ is to note that with positive values of δ, fairly large values of σ are necessary to generate realistic levels of capital in the model. But patient people with negative δ's will accumulate substantial wealth even with low intertemporal elasticities of substitution. Note that this dispute about calibrating the model assumes that national wealth is attributable entirely to lifecycle saving. If a significant share of the capital stock arises from precautionary saving or bequests, however, one cannot use aggregate wealth to guide the choice of these two parameters.

Auerbach and Kotlikoff (1987) develop an elaborate simulation model to study the dynamic effects of tax and fiscal policy. They confirm Summers' central conclusion that a shift from income taxation to consumption taxation would significantly increase the capital stock, but they do not report interest elasticities of saving.[42] Auerbach and Kotlikoff use a time preference rate of 0.015 in their simulations. They argue that this rate produces a reasonable level of capital in their model (but see the previous paragraph) and that negative rates of time preference would produce "highly unrealistic" growth rates of consumption with age (p. 70).

Gravelle (1994) carefully reviews the studies discussed here as well as other research that analyzes the effect of reducing the taxation of capital income. She notes that the saving responses in most of these studies are very large even for moderate values of the intertemporal elasticity of substitution. For example, Auerbach and Kotlikoff assume that σ equals 0.25 and report that an elimination of capital taxation would more than double the saving rate for almost 10 years. Gravelle argues that "the implausibility of this result suggests that either some revision

[42] It is difficult to calculate elasticities from their reported results because saving is increased not only by an increase in the after-tax rate of return, but also by a change in the timing of taxes (people pay consumption taxes later in their lives than they pay income taxes). Interestingly, Auerbach and Kotlikoff claim that the intertemporal elasticity of substitution has little effect on the percentage increase in the steady-state capital stock because it has roughly proportional effects on the initial amount of capital and the predicted change in capital.

in the model or use of a much lower" σ would be appropriate (p. 42). She does not propose a specific alternative model, however.

5. The Effect of Interest-Rate Changes on Bequest Leavers

The previous sections of the paper explore the effect of interest-rate changes on the consumption and saving of people who behave in accordance with the lifecycle model. Although this model is economists' most familiar representation of consumer behavior, it clearly does not describe all people. One alternative model asserts that individuals' planning horizons extend beyond their lifetimes so that their saving decisions reflect concern for their children as well as for themselves. This section of the paper examines the effect of changes in interest rates on consumption and saving that is influenced by a desire to leave bequests to future generations.

To begin, note that the basic lifecycle model analyzed in the earlier sections has no role for bequests. Individuals care only about their own consumption, so they would like to consume all of their wealth during their lifetime, and they have no uncertainty about the length of their lives, so there is no reason why they would be frustrated in this desire. In the real world, however, bequests clearly exist, and they probably comprise a significant portion of aggregate wealth. Kotlikoff and Summers (1981) estimate that intergenerational transfers represent over half of total wealth in the United States. Modigliani (1988) argues that the true number is probably closer to one-fourth, but that is still a substantial share.[43] Thus, understanding how changes in interest rates affect the saving of bequest-leavers is clearly important to understanding how interest-rate changes affect aggregate saving.[44]

---

[43] See also Kotlikoff (1988) and Gale and Scholz (1994).

[44] Bequests may be important for aggregate saving even if most people do not leave significant bequests. Avery and Kennickell's (1991) exploration of the Survey of Consumer Finances shows that among people over age 70, median saving is less than zero (the average person is dis-saving) but mean saving is greater than zero (the group as a whole is still accumulating wealth).

Unfortunately, the effect of interest-rate changes on saving for bequests depends entirely on the motivation for those bequests. Disparate motivations exist, and economists have addressed three main possibilities. First, people may leave bequests accidentally, because they die earlier than they expected. Second, people may leave bequests of a fixed size, perhaps because they value the giving itself or to influence their children's behavior. And third, people may leave bequests altruistically, because their children's well-being enters their own utility functions. This last motivation for leaving bequests is probably the most important, although there is little direct evidence on this point. The remainder of this section takes up each motivation and its effect on the interest elasticity of saving in turn. I conclude that the saving of altruistic bequest leavers responds positively and quite strongly to increases in interest rates.

Accidental Bequests

The lifecycle model used in the previous sections of this paper assumes a world of certainty. In the real world, prudent people who are uncertain about their lifespan will accumulate sufficient resources to support themselves in retirement for many years. If these people die unexpectedly early, their remaining financial assets will be passed on to their heirs.[45]

The effect on the interest elasticity of saving of people's uncertainty about their lifespans is examined by Engen (1992) and reviewed at the end of the second section of this paper. As explained there, this uncertainty definitely reduces the elasticity, but the magnitude of the reduction is unclear.

Fixed Bequests

Some people aim to leave a bequest of a fixed amount. They may want to give their children

---

[45] Auerbach, Kotlikoff and Weil (1992) show that an increased share of the elderly's wealth is annuitized--meaning that it represents a claim to a stream of payments that ends with a person's death--compared to several decades ago. They note that increased annuitization should reduce the extent of accidental bequests.

or a charitable organization a gift of a certain size, or they may feel that the promise of an inheritance wins them greater attention from their children, as suggested by Bernheim, Shleifer and Summers (1985).

This type of bequest is just consumption at the end of life, so it pushes planned consumption further into the future than for a pure lifecycle individual. As a result, it increases the size of the income effect relative to the calculations in the previous section, and reduces the overall interest elasticity of saving.[46] The magnitude of the reduction depends on the size of the intended bequest.

Altruistic Bequests

Most people bear children, and then care about their children's well-being. Although their children may be better off economically than they are (due to productivity gains over time), these people may care enough about their children to leave them bequests. Such altruistic bequests have received the most attention from economists, perhaps because their implications are so different from the implications of the basic lifecycle model. These differences arise because people leaving altruistic bequests behave effectively as if they were infinitely-lived.[47]

What is the interest elasticity of saving for infinitely-lived agents? They will save whenever the (after-tax) interest rate exceeds their rate of time preference, and dis-save whenever the (after-tax) interest rate falls below their rate of time preference. Saving increases the capital stock, bringing down the marginal product of capital and the interest rate; dis-saving has the opposite effects. Thus, a steady state is achieved when the after-tax interest rate equals the rate of time preference: $(1-x)\ r = \delta$, where x equals the tax rate on capital income. This condition

[46] Moreover, the absolute value of the elasticity is decreased because initial saving is larger. See Evans (1983).

[47] Barro (1974) shows that if people care about their children, and the children care about *their* children, and so on, then the infinite stream of generations becomes effectively an infinitely-lived dynasty.

implies that capital supply is infinitely elastic in the long run. Of course, the short-run elasticity of saving is finite, and its magnitude depends on other features of preferences and the economy. This issue of how quickly the economy moves through the transition to the long run is beyond the scope of this paper. Nevertheless, it seems clear that altruistic bequest-leavers have a substantially positive interest elasticity of saving.

6. The Effect of Interest-Rate Changes on Target Savers

The lifecycle model presumes that people choose their current consumption and saving in order to maximize their lifetime well-being given their lifetime resources. The assumption that many people are forward-looking enough to save for their future needs seems quite reasonable, but it is at least possible that some forward-looking people do not determine the amount of their saving in a specific optimizing framework. Instead, these people may simply save enough early in their lives to achieve a given level of consumption later. Such behavior is often called "target saving." This section of the paper examines the effect of changes in interest rates on the consumption and saving of target savers.

Some people may do target saving as a substitute for lifecycle saving, as suggested in the previous paragraph. A common financial-planning technique, for example, is to save enough when working to attain a chosen fraction of pre-retirement income in retirement. Other people may do target saving as a complement to lifecycle retirement saving. People who save for a vacation, for a downpayment on a house, or for a child's education, may all be target savers. The crucial characteristic of target savers is that their future consumption is not sensitive to interest rates in the way that the future consumption of lifecycle savers is. Calculations presented below show that this difference has an important effect on the interest elasticity of saving.

The first part of this section examines the interest elasticity of simple target saving, and the second part considers the interest elasticity of saving by people who follow the advice of financial planners. Both elasticities are generally negative.

Simple Target Saving

A simple form of target saving involves saving just enough today (or over a number of years) to reach some desired level of assets at a specific date in the future. A classic example is saving for a child's college education. Not surprisingly, an increase in the interest rate reduces the amount of saving needed each year in order to reach that target amount. The interest elasticity of this simple target saving is negative, as shown formally in Appendix B.

This simple target saving undoubtedly describes some saving by some people. But more complicated forms of target saving may be important, too. For example, in contrast to the example of one target at a single date in the future, described in the previous paragraph, people who are saving for retirement want to reach a set of targets over a number of years. A more fundamental issue that is neglected by models of simple target saving is the effect of wealth revaluations: if a decline in the interest rate increases someone's financial wealth, would he or she reduce saving for the target, raise the target, or apply the greater wealth to other future needs? To answer this question, one needs a more complete theory of how a target saver saves over a lifetime. The next part of this section considers a complete model of this sort.

The Financial-Planning Model

Determining the optimal level of consumption and saving in a lifecycle model is a complex problem, as the lifecycle formulas in Appendix A suggest. To cope with this complexity, many people seek the advice of professional financial planners, either in person, through books, or via computer software. For example, several large mutual-fund companies now offer their customers low-cost software that asks the user a set of questions about his or her financial situation and goals, and then calculates an appropriate level of saving for the person.

The specific advice offered by these planners varies, of course. But one common recommendation is that people should decide what fraction of pre-retirement consumption they want to achieve during retirement, and then save enough to reach that goal.[48] This fraction is called the "replacement rate." The remainder of this section considers the interest elasticity of saving for people following this "financial-planning model."[49]

<u>Theory</u>. The financial-planning model is actually quite similar to the lifecycle model, because consumption and saving behavior are forward-looking in both cases. Thus, both models propose a path of future consumption that will exhaust all lifetime resources by the end of one's life (that is, a path of consumption which satisfies the intertemporal budget constraint).

The crucial difference between the financial-planning model and the lifecycle model is that the slope of the lifecycle consumption path generally depends on the interest rate, while the slope of the financial-planning consumption path does not. Specifically, the slope of a lifecycle consumption path equals $\sigma(r-\delta)$, where σ is the intertemporal elasticity of substitution and δ is the rate of time preference. For σ greater than zero, the consumption path is upward-sloping as long as r is greater than δ (as in all of the numerical examples in the fourth section of this paper). In contrast, the financial-planning consumption path is flat until retirement, and then is flat again at a lower level after retirement.[50] Thus, the interest rate does not affect the slope of the consumption path in the financial-planning model.

What do this similarity and difference tell us about the interest elasticity of saving in the financial-planning model? Recall that the effect of interest-rate changes on lifecycle consumers can be decomposed into three pieces: a substitution effect, an income effect, and a wealth effect.

---

[48] For example, see the articles in Doyle and Johnson (1991). At least some financial-planning tools ask what fraction of pre-retirement *income* one wants to achieve during retirement, although it is more natural to think about the share of consumption. The following analysis assumes that people choose a replacement rate for consumption.

[49] I am grateful to Bill Gale for suggesting that I address this issue.

[50] The lifecycle model can be extended by allowing tastes to change at retirement. This could make consumption late in life lower than consumption while working, but it would leave the slope of the consumption path during each part of life a function of the interest rate.

The substitution effect for lifecyclers is based on a change in the slope of the consumption path caused by the change in the relative price of consumption at different times. Because interest-rate changes do not affect this slope under the financial-planning model, the substitution effect in that model is zero. The income and wealth effects are based on changes in the intertemporal budget constraint. Because this constraint applies to financial-planning consumers as well, their consumption and saving will also be subject to income and wealth effects. The magnitude of these effects is generally different than for lifecyclers, however, because the timing of consumption over the lifetime is generally different.

The income effect captures the fact that, with higher interest rates, fewer current dollars are needed to fund a given amount of future consumption. Thus, this part of the response is similar to the effect of interest-rate changes in the simple target-saving framework described earlier in this section. The income effect works in the direction of a negative interest elasticity of saving.

The wealth effect captures the effect of higher interest rates on the present discounted value of people's expected future income. Higher interest rates reduce this value, which encourages people to save more and consume less. Thus, the wealth effect works in the direction of a positive interest elasticity of saving.

Numerical formulas for the strength of these two effects are derived in Appendix B. Because the effects operate in opposing directions, one cannot determine theoretically whether the interest elasticity of saving in the financial-planning model is positive or negative. But simple calculations reveal that the elasticity is almost always negative, as I show next.

Illustrative calculations. Tables 9 and 10 use the formulas derived in Appendix B to quantify the response of financial-planning savers to a change in the interest rate. The calculations are based on the same assumptions that underlie the lifecycle calculations in the fourth section of the paper, with the following exceptions. First, there is no need to specify an intertemporal elasticity of substitution or rate of time preference, because people following the financial-planning model do not choose a consumption path based on explicit utility maximization. Instead, these people choose a replacement rate for consumption, which I specify below. Second, I assume that labor earnings do not increase over a person's lifetime. This assumption simplifies the problem and is consistent with at least some financial-planning advice.

Table 9 presents the results of the calculations for a replacement rate equal to 80 percent. The interest elasticities of saving for a 30-year old and a 50-year old are negative, regardless of the type of financial wealth they hold. In both cases, the income effect outweighs the human wealth effect, and the effect of financial-wealth revaluation is too small (even if non-zero) to make much of a difference. The 70-year old experiences a much smaller income effect, however, because she has fewer years of consumption left, and a larger financial wealth effect because she has accumulated more wealth. As a result, a 70-year old who holds her wealth in the form of physical capital has a positive interest elasticity of saving.

These results are not sensitive to the individual's chosen replacement rate. Table 2 shows the interest elasticity of saving under the lifecycle model with an intertemporal elasticity of substitution equal to zero. As explained below, this situation is equivalent to the financial-planning model with a replacement rate of 100 percent. Yet, despite this difference in replacement rates, the values reported in Tables 2 and 9 are quite similar.

Table 10 presents some calculations of the aggregate interest elasticity of saving corresponding to the parameters used in Table 9. The aggregation assumptions are the same as those used for Table 8 based on the lifecycle model and discussed in the fourth section of the paper. For a combined productivity and population growth rate of 3 percent, and declining physical capital holding over a lifetime, the elasticity is -1.17.

<u>Evaluation</u>. As noted above, the key difference between the lifecycle and financial-planning models is that the interest rate does not affect the slope of the consumption path in the financial-planning model. In other words, the model implies that people do not respond to a change in the price of future consumption relative to current consumption (although they do respond to the

related change in their lifetime well-being).

How plausible is this feature of the model? If someone has chosen a replacement rate for consumption, then changes in relative prices are indeed irrelevant. Yet, the financial-planning model does not specify how people should choose (or do choose) a replacement rate. This indeterminacy leaves open the possibility that changes in interest rates could affect the desired replacement rate.

For example, consider someone choosing between saving enough to replace 80 percent of consumption when retired, and saving enough to replace 90 percent. Less *additional* saving is needed to reach the higher replacement rate when the interest rate is higher. This makes the higher replacement rate relatively more appealing, which might make the person more likely to choose it. If the person does raise his or her chosen replacement rate in response to higher interest rates, then interest rates do affect the slope of the consumption path over the lifecycle. Thus, this example shows how a person who does not solve the formal lifecycle model may still respond to relative prices in the way that the lifecycle model predicts.

One piece of evidence that people think in this way comes from estimates of the intertemporal elasticity of substitution, which are reviewed in the third section of the paper. Note that a financial-planning consumer who chooses a 100 percent replacement rate is indistinguishable from a lifecycle consumer with an elasticity of substitution equal to zero: both maintain constant consumption over their entire lives regardless of the interest rate. Therefore, if the financial-planning model applied to most people, one would expect the estimated intertemporal elasticity of substitution to be zero. In fact, the estimates--though not without problems--imply an elasticity around 0.33. This suggests that most people do respond to a change in the relative price of current and future consumption.

7. The Effect of Interest-Rate Changes on Rule-of-Thumb Consumers

The lifecycle model and financial-planning model assume that people choose their current consumption and saving by looking forward to their income and desired consumption over their entire lifetime. The model of altruistic bequests applies to people who look further forward, into their children's lives. But many people appear not to be as far-sighted as these models predict. Thus, an alternative model of behavior asserts that individuals have short planning horizons and pick their consumption and saving based primarily on their current income. This section of the paper examines the effect of changes in interest rates on these people's consumption and saving.

There is substantial evidence that many households do not smooth consumption over time as predicted by the lifecycle model. Some of this evidence comes from data on aggregate income and consumption: Campbell and Mankiw (1989) examine the relationship between these series and conclude that roughly half of aggregate consumption is done by individuals who follow the "rule of thumb" of consuming their entire current income. Similar evidence arises from data on the income and consumption of individual households. For example, Zeldes (1989) finds that borrowing constraints affect a significant portion of the population, and Carroll and Summers (1991) show that the pattern of consumption over a person's lifetime often matches the pattern of his or her income quite closely. Yet further confirmation comes from households' responses to changes in taxes and government benefits: see Poterba (1988), Wilcox (1989), and Shapiro and Slemrod (1994), for example.[51]

The common feature of this evidence is that many individuals' consumption is more sensitive to their current income than forward-looking models imply. But it is less clear exactly how these people determine their consumption and saving, which makes it difficult to know how they respond to a change in interest rates. The first part of this section examines the behavior of people who do not smooth consumption, focusing on rule-of-thumb consumers who set consumption equal to their cash flow. The following part explores the effect of interest-rate changes on household cash flow; I conclude that a rise in interest rates increases aggregate

---

[51] The empirical evidence on consumption smoothing is reviewed by Deaton (1992) and Browning and Lusardi (1994).

household cash flow, but may lower the cash flow of rule-of-thumb individuals.

Rule-of-Thumb Behavior and Interest Rates

There are at least two types of people who do not smooth consumption over time. One type is people who would like to smooth their consumption, but face binding borrowing constraints. These people are thinking in lifecycle terms, but their desired financial wealth is negative, which is generally not feasible. This situation is more common among younger people with a higher rate of time preference or a steeper age-earnings profile (as shown in the fourth section of the paper). In the simplest case, liquidity-constrained consumers have no assets and set consumption equal to income. If credit markets allow limited borrowing, these consumers may have some debts, but would still set consumption equal to income.[52]

Other people who do not smooth consumption simply may not think about their behavior in a systematic forward-looking way. For example, Shefrin and Thaler (1988) and Thaler (1994) argue that people have a set of "mental accounts" with different propensities to consume out of different types of income and wealth. Some myopic individuals probably consume all of their income year-by-year, or month-by-month. Others might save a fixed fraction of their income, or save the excess of their income over some slowly-changing level of consumption, or consume a portion of current income plus a portion of any financial wealth they hold. These people's wealth could be negative or positive, but it is probably fairly small either way.

For both types of people with short horizons, consumption depends on contemporaneous income and perhaps a small amount of (positive or negative) financial wealth.[53] Therefore,

---

[52] These sentences assume that the liquidity constraint is binding in each period. If a constraint is not currently binding but is expected to bind at a future time, the consumer may have positive financial wealth and set consumption not according to the unrestricted lifecycle model nor simply equal to income. See Deaton (1992) and Browning and Lusardi (1994) for further discussion of this point.

[53] This is just a traditional Keynesian consumption function.

changes in interest rates can affect these people's behavior by changing either their income or their wealth. Earlier sections of the paper discussed financial-wealth revaluations at length, so I focus here on the effect of interest-rate changes on income. For simplicity, I consider only rule-of-thumb consumers who set their consumption equal to their income.

More specifically, I assume that these people set consumption equal to their *cash flow*, where cash flow is defined as income less interest payments. Disposable (after-tax) income in the National Income and Product Accounts (NIPAs) is divided not between consumer expenditures and saving, as one might expect, but between expenditures, saving, and interest payments.[54] Yet, interest payments are conceptually quite similar to interest receipts, which are included in disposable income, because both flows represent the servicing of stocks of wealth: interest receipts correspond to positive elements of wealth (assets) and interest payments correspond to negative elements (liabilities). The concept of cash flow used here treats the receipt and payment of interest symmetrically.[55]

How would a change in interest rates affect household cash flow? In the short run, households' stocks of assets and liabilities can be considered fixed, so interest-rate changes affect cash flow only to the extent that interest rates paid or received are adjustable over time rather than set in advance. For example, households with fixed-rate mortgages experience no change in their cash flow when interest rates rise, while households with adjustable-rate mortgages are hurt. A change in the interest rate tends to move both the interest received by households and the interest paid in the same direction; the crucial empirical question of which effect dominates is considered in the next part of this section. Over time, the change in cash flow depends on both the change in interest rates and any adjustment in the amount of assets and liabilities held. As in the earlier discussion of lifecycle consumers, I restrict my attention to the short-run response.

---

[54] There is also a much smaller category of transfer payments to foreigners.

[55] NIPA income includes a number of imputations, so the actual flow of money into households differs from the measure of cash flow used here.

Suppose that an increase in interest rates raises a rule-of-thumb consumer's interest payments and thus reduces his or her cash flow. This person reduces consumption by the same amount, leaving his or her saving unchanged at zero. At the same time, *aggregate* saving must have increased, because aggregate income is unchanged and aggregate consumption has declined because of this person's action.[56] These observations about aggregate saving and the saving of the rule-of-thumb consumer can be reconciled by noting that this individual's increased interest payments represent higher interest receipts for someone else. Suppose (for simplicity) that this other person is a lifecycle consumer. Changes in her current capital income have no effect on her consumption, which adjusts only to changes in her lifetime resources. Thus, the increase in current income flows directly through to an increase in her saving, which is in addition to any change in her saving for the lifecycle reasons analyzed earlier. This analysis implies that any change in the consumption of rule-of-thumb consumers produces a corresponding change in aggregate saving.

The Effect of Interest-Rate Changes on Household Cash Flow

The preceding discussion shows that changes in interest rates can affect aggregate consumption and saving by changing the cash flow of rule-of-thumb consumers. Unfortunately, the direction of this effect is theoretically ambiguous: a rise in interest rates, for example, increases the interest paid by households but also the interest received. It is an empirical question whether household cash flow is positively or negatively correlated with interest rates, and therefore whether rule-of-thumb consumers push the interest elasticity of saving in a positive or negative direction. The remainder of this section tries to answer this question in several steps.[57]

The first step is to determine whether the interest receipts of the household sector are

---

[56] Aggregate income is unchanged because the country's productive capacity is essentially fixed in the short run, and I abstract from business cycles (as in the second section of the paper).

[57] The magnitude of the effect is an important empirical question as well, but I focus on the easier question of sign.

larger or smaller than its interest payments. The first panel in Figure 2 shows the NIPA data on household interest receipts and payments as a share of cash flow.[58] These series are misleading, however, because they include interest income that is actually received by pension funds and insurance companies, and they ignore mortgage interest expense, which is assumed to be paid by fictional owner-occupied housing entities to which households pay imputed rent. The second panel of Figure 2 adjusts for these problems, which brings household receipts and payments much closer together than they appear in the raw data. Even these adjusted values are not quite comparable, however, because the tax treatment of different interest flows is different: most interest income is taxable, and although mortgage interest is tax-deductible, other consumer interest is not. The third panel of Figure 2 makes a crude adjustment for these tax differences, showing that households' after-tax interest receipts narrowly exceeded their interest payments for most years between 1967 and 1990. The gap is generally rather small, however, and has actually reversed during the past several years.[59]

The second step in determining the effect of interest-rate changes on household cash flow is ascertaining whether households' interest receipts adjust more or less quickly than their interest payments to changes in market interest rates. Goodman, Luckett and Wilcox (1988) suggest a simple procedure for estimating these adjustment speeds, which I implement in Appendix C and summarize here. I construct implicit average yields for interest receipts and interest payments on an annual basis, and regress those yields on the contemporaneous and three lagged values of the one-year Treasury bill rate. For interest receipts, the first two estimated coefficients sum to 0.53, while all four coefficients sum to 0.87. For mortgage interest payments the corresponding values are 0.31 and 0.77, and for other interest payments they are 0.16 and 0.24. These results

[58] Appendix C contains details on the construction of these series and those discussed below.

[59] The recent reversal of the historical relationship is attributable primarily to a significant runup in household debt. Both the surge in interest receipts (as a share of cash flow) in the early 1980s and their subsequent decline can be explained almost entirely by the swing in market interest rates and its effect on the average yield on household interest-bearing assets. On the other hand, the same swing in market rates would imply a larger decline in interest payments (as a share of cash flow) than actually occurred. The difference is an increase in the stock of household liabilities.

show that the implicit interest rate on interest receipts adjusts much more quickly than the implicit interest rates on interest payments.

Two notes of caution are in order, however. First, the relationships in question may not be stable over time. A shift toward adjustable-rate mortgages, for example, will increase the responsiveness of interest payments to market interest rates. Second, the relationships may not be symmetric for increases and decreases in market interest rates. A decline in rates will spur a round of mortgage refinancing (which lowers the average interest rate on payments) that is not repeated when rates rise. Further exploration of these topics is beyond the scope of this paper, so I take the simple results at face value.

To review, the household sector currently receives slightly less interest than it pays out, but its receipts respond much more quickly than its payments to changes in market interest rates. The combined effect is that an increase in the interest rate generally raises households' interest receipts more than their interest payments, especially during the first few years. As a result, household cash flow increases, consumption rises, and aggregate saving falls. It appears that rule-of-thumb behavior contributes to a negative interest elasticity of aggregate saving.[60]

This conclusion could be incorrect, however, if rule-of-thumb consumers pay out more interest than they receive, while lifecycle consumers and bequest leavers receive more interest than they pay out. This distribution of interest receipts and payments is consistent with the small difference between receipts and payments for the household sector as a whole, and also with the earlier description of rule-of-thumb people as having low or negative wealth: if a household's liabilities exceed its assets, one would not be surprised if its interest outgo exceed its interest income. If receipts and payments are distributed in this way, then rule-of-thumb consumers' cash flow might be hurt by an increase in interest rates despite the more rapid adjustment of receipts.

---

[60] Cantor (1989) and Goodman, Luckett and Wilcox (1988) also conclude that households are likely to experience positive changes in cash flow when interest rates rise. These papers carefully examine the responsiveness of different categories of interest payments and receipts to changes in market interest rates.

As a result, their consumption would decline, and they would contribute to a positive interest elasticity of aggregate saving, not a negative one. It is not clear, though, whether rule-of-thumb consumers actually *are* hurt by an increase in the interest rate. The issue can be investigated using household-level data on assets and liabilities, if one makes assumptions about which households are likely to follow rule-of-thumb behavior (such as those with low or negative wealth). Such an investigation lies beyond the scope of this paper, however.[61]

8. Other Evidence on the Interest Elasticity of Saving

The introduction to this paper argues that the most compelling approach for determining the interest elasticity of saving is an indirect one, in which models of people's behavior are combined with estimates of certain features of their preferences. Thus, the second through fourth sections of the paper examine the effect of interest-rate changes on lifecycle consumers, while the fifth through seventh sections cover bequest leavers, target savers, and rule-of-thumb consumers. Although this indirect approach generates useful information about the interest elasticity of saving, it is by no means a panacea. Economists do not know how much of aggregate saving can be described by each of these models, and even in the context of a particular model, it is difficult to determine the response of saving to interest-rate changes. Therefore, this penultimate section of the paper briefly reviews some other types of evidence on the interest elasticity of saving. I begin with estimates of aggregate consumption functions, then turn to analyses based on differences in tax rates, and conclude by reviewing the U.S. experience of the 1980s.

Estimates of Aggregate Consumption Functions

[61] Goodman et al point out that most debt-holders have substantial assets as well. But these people may be lifecyclers or altruistic bequest leavers, whose response to interest-rate changes is considered in earlier sections. The question here is whether people whose consumption responds primarily to their current cash flow have substantial assets to offset their debts, and that seems less likely.

One direct approach to determining the interest elasticity of saving is to estimate an aggregate consumption function that includes the interest rate as an explanatory variable. An aggregate consumption function links total national consumption to some of the factors that influence it, such as national income, national wealth, and so on. When the interest rate is included as one of those factors, its coefficient can be interpreted as the effect of interest-rate changes on aggregate consumption and thus on aggregate saving.

In the seminal study of this type, Boskin (1978) estimates an interest elasticity of saving equal to 0.4. Unfortunately, later work by Howrey and Hymans (1978), Friend and Hasbrouck (1983), and others shows that this estimate is quite sensitive to the measurement of the variables (especially the rate of return) and to the sample period used in the estimation. For certain variations in the specification the estimated elasticity is negative, and for many variations it is less than 0.4.[62] On balance, the estimation of aggregate consumption functions has provided little reliable evidence regarding the interest elasticity of saving.

The unreliability of these estimates is not very surprising, given the fundamental problems with this approach.[63] Some of these problems are shared by the Euler-equation estimates of the intertemporal elasticity of substitution reviewed earlier. For example, it is hard to measure the expected real return on saving, because of difficulties in modeling both inflation expectations and expected nominal returns on equity. Other problems that are prominent in estimates of aggregate consumption functions are largely avoided by Euler-equation estimates. For example, aggregate consumption functions interpret a correlation between the interest rate and consumption as the causal effect of the interest rate. It is just as likely, however, that the correlation reflects the effect of consumption (and thus saving) on the interest rate. Euler equations estimated with household data avoid this serious endogeneity problem. Another problem with aggregate consumption functions is modelling expected future income; the Euler-equation approach does not require such a model. Finally, the estimated consumption functions often include financial

---

[62] See Gravelle (1994, page 27) for a review.

[63] For further discussion of these problems and others, see Summers (1982, 1984).

wealth as an additional explanatory variable. This means that the effect of interest-rate changes on consumption via the revaluation of financial wealth appears in the coefficient on wealth and not in the coefficient on the interest rate itself.

Tax-Based Evidence

Capital income taxes have a significant effect on the rate of return earned on saving by most households. Therefore, differences in tax rates across households or across time can be used to study the responsiveness of saving to changes in the rate of return.

One source of variation in the after-tax return to saving is provided by Individual Retirement Accounts (IRAs), 401(k) plans, and other tax-favored saving arrangements. Unfortunately, there is no consensus about the effect of these tax incentives on saving, and it is beyond the scope of this paper to evaluate the extensive literature on this topic. For evidence that IRAs and 401(k) plans raise saving, see Venti and Wise (1990) and Poterba, Venti and Wise (1995) among others; for opposing evidence, see Gale and Scholz (1994) and Engen, Gale and Scholz (1994) among others.

The link between the interest elasticity of saving and the response of saving to targeted incentives is fairly weak in any case. Suppose, for example, that IRAs do increase saving significantly. This effect could arise because the interest elasticity of saving is large, or because people respond to advertising about IRAs and to the up-front deductibility of IRA contributions, as emphasized in so-called "behavioral" theories of saving.[64] On the other hand, if IRAs do not increase saving, that might be because the interest elasticity of saving is small, or because the legal limit on IRA contributions means that many taxpayers receive no higher return on an extra dollar of saving.

---

[64] For a description of these theories and their relevance to national saving policy, see Bernheim (1996).

Another source of variation in the after-tax rate of return is changes in the tax law over time. A prominent example would be the Tax Reform Act of 1986 (TRA86), except that the many provisions of that law had opposing impacts on the effective tax rate. Thus, Feenberg and Skinner (1990, page 72) argue that TRA86 increased the after-tax return to saving for most people, while Hendershott (1990, page 83) concludes that the opposite is true. Clearly, it is impossible to determine the interest elasticity of saving from a tax change whose effect on the after-tax rate of interest is uncertain. Moreover, there were many other forces affecting aggregate saving during the mid-1980s, and it is not clear how one could isolate the effect of the tax law.

The Experience of the 1980s

Several developments in the U.S. economy during the 1980s would be expected to increase household saving. First, the rate of return to saving was higher in the 1980s than in previous decades, because of a combination of higher real interest rates, lower tax rates, and increased access by households to market rates of return. Second, the distribution of income became less equal during the 1980s, which presumably shifted income toward people with higher saving rates. Still, the personal saving rate was lower during the 1980s than in previous decades.[65] This experience appears to show that the interest elasticity of saving is negative or zero, but such an interpretation seems unwarranted for at least two reasons.

First, interest rates may have been high in the 1980s partly *because* household saving was low. The interest rate is an endogenous variable that equilibrates the supply and demand for loanable funds. If the observed rise in the interest rate was caused solely by an increase in the demand for funds, then the change in household saving would indeed reflect the slope of the supply curve. But if the rise in the interest rate was caused by an exogenous decline in the supply of funds, then the correlation between interest rates and saving would shed no light on the interest elasticity of saving. Unfortunately, the relative importance of demand shifts and

[65] The magnitude of the decline depends on exactly how saving is measured, but all measures show a decline of some degree. See Congressional Budget Office (1993).

supply shifts in explaining the increase in interest rates in the 1980s is not known. The increase in the federal budget deficit clearly boosted the demand for loanable funds, but other factors simultaneously diminished the supply, as noted in the next paragraph.

Second, several developments in the U.S. economy during the 1980s would be expected to reduce household saving for a given level of interest rates.[66] One example is the surge in household wealth, including large gains in stock and bond prices, and an erosion in the real value of mortgages caused by the high inflation of the late 1970s and early 1980s. Another example is the greater ease of borrowing, caused by the expansion of home equity loans and by broader access to credit cards with generous terms.[67]

9. Conclusion

This paper reviews economists' knowledge about the interest elasticity of saving. Direct estimates of this elasticity suffer from a number of serious problems that make the estimates quite sensitive to small changes in specification. As a result, this paper focuses on an indirect approach to determining the elasticity that combines models of individual behavior with estimates of certain features of individuals' preferences. The paper considers several alternative models of behavior, examining the effect of interest-rate changes on the consumption and saving of people who follow the lifecycle model, people who plan to leave bequests, people who save to reach a fixed target, and people with very short planning horizons. The paper also reviews the evidence regarding the critical preference parameters. Almost all of the analysis is restricted to the behavior of households in the short run; the response of businesses to changes in interest rates, and the long-run response of households, are topics for another paper.

---

[66] For further discussion of these and other factors, see Congressional Budget Office (1993).

[67] For an analysis of the effect of second mortgages on household saving, see Manchester and Poterba (1989).

My conclusion regarding the interest elasticity of saving is two-fold. First, economists' understanding of the response of household saving and consumption to changes in interest rates is quite limited, despite a large volume of research on the topic. Alternative models of consumer behavior imply different magnitudes for the interest elasticity of saving, and even different signs. Each model probably describes the behavior of some people, and it is not clear which model best characterizes the behavior of the "average saver." Thus, it is simply not possible to provide a precise estimate of the interest elasticity of saving with any confidence.

Despite the uncertainty, the models that likely describe the behavior of the people who account for most of aggregate saving imply positive interest elasticities. Thus, the paper's second conclusion is that the short-run interest elasticity of saving is probably positive. A basic lifecycle model with empirically-supported parameters easily generates an elasticity around 0.5, although the magnitude is sensitive to the exact parameter choices. Adding uncertainty to the basic model reduces the absolute value of the elasticity, perhaps significantly, but it does not transform positive elasticities into negative ones. Adding altruistic bequests to the lifecycle model increases the elasticity, although other, probably less important, motives for leaving bequests have the opposite effect. Target saving--including saving behavior suggested by many financial planners--implies a significant negative interest elasticity of saving, but there is at least some evidence that many people respond to interest rates in a way not predicted by the target-saving model. Rule-of-thumb consumers may contribute to a positive or negative interest elasticity of saving, depending on the types of assets and liabilities they hold. It is unlikely, however, that these people do a large enough share of aggregate consumption or saving to have a substantial effect on the aggregate elasticity. In sum, the combination of theory and empirical evidence suggests on balance that the aggregate interest elasticity of saving is unlikely to be negative, and may be substantially positive.

Appendix A: Numerical Formulas for Lifecycle Consumers

The second section of the paper presents the theory of how lifecycle consumers respond to a change in the interest rate. As explained in that section, the total effect of an interest-rate change on consumption and saving can be decomposed into a substitution effect, an income effect, and a wealth effect. This appendix derives numerical formulas for the strength of these three effects.

Suppose that an individual's well-being over her future life can be described by:

$$U = \int_{\tau}^{T_2} \frac{1}{1-\frac{1}{\sigma}} C_t^{\left(1-\frac{1}{\sigma}\right)} e^{-\delta(t-\tau)} dt \quad ,$$

where U is utility, C_t is consumption in period t, σ is the intertemporal elasticity of substitution, δ is the rate of time preference, and τ is the person's current age and T_2 the person's age at death (both measured from the beginning of her working life).[68] Also suppose that an individual's intertemporal budget constraint can be written as:

$$\int_{\tau}^{T_2} C_t e^{-r(t-\tau)} dt = \int_{\tau}^{T_2} Y_t e^{-r(t-\tau)} dt + F_{\tau} \quad ,$$

where Y_t is the sum of labor and pension income in period t, r is the interest rate, and F_{τ} is financial wealth at time τ.[69] Denoting the present value of all income (the right side of the

[68] This utility function displays constant relative risk aversion, with the coefficient of relative risk aversion equal to the inverse of the intertemporal elasticity of substitution ($1/\sigma$). Economists disagree about whether these two concepts should be linked so closely, and some recent work has broken the link at the cost of introducing less tractable utility functions. For further discussion, see Browning and Lusardi (1994).

[69] More accurately, Y is all income except that derived from existing financial wealth or new saving. Thus, in addition to labor earnings and pension receipts, it includes transfer payments, gifts received, and similar items.

equation) as W_τ, the budget constraint becomes:

$$W_\tau = \int_\tau^{T_2} C_t \, e^{-r(t-\tau)} \, dt \ .$$

Combining the budget constraint and the utility function, the first-order condition for utility maximization is:

$$C_t = C_\tau \, e^{\sigma(r-\delta)(t-\tau)} \ .$$

Substituting the first-order condition into the budget constraint yields the following formula for optimal consumption:

$$C_\tau = \mu_\tau \, W_\tau = \frac{\theta}{1 - e^{-\theta(T_2-\tau)}} \, W_\tau \ ,$$

where μ_τ is the marginal propensity to consume out of wealth, and θ equals r - $\sigma(r\text{-}\delta)$. The derivative of C_τ with respect to the interest rate is:

$$\frac{\partial C_\tau}{\partial r} = W_\tau \, \frac{\partial \mu_\tau}{\partial r} + \mu_\tau \, \frac{\partial W_\tau}{\partial r} \ .$$

The first term of this expression represents the substitution and income effects, and the second term represents the wealth effect. I now discuss the effects in the corresponding order.

The Substitution and Income Effects

The first term in the previous equation can be rewritten as:

$$W_\tau \, \phi_\tau \, \frac{\partial \theta}{\partial r} \ ,$$

where

$$\phi_\tau = \frac{1 - \left(1 + \theta(T_2 - \tau)\right) e^{-\theta(T_2 - \tau)}}{\left(1 - e^{-\theta(T_2 - \tau)}\right)^2} \quad .$$

Consider the term $\partial\theta/\partial r$, where θ equals $r - \sigma(r-\delta)$, as above. The substitution effect involves a change in the slope of the consumption path holding the present value of planned future consumption constant, while the income effect is based on a change in that present value. In other words, the substitution effect corresponds to the second appearance of r in the formula for θ (which comes from the first-order condition), and the income effect corresponds to the first appearance of r in the formula for θ (which comes from the budget constraint). Therefore, when we evaluate the expression above, we can calculate the substitution effect by *holding constant the first r in the formula for θ* and the income effect by *holding constant the second r in the formula for θ.* As a result, the substitution effect can be written as:

$$-\sigma \ \phi_\tau \ W_\tau \ ,$$

and the income effect can be written as:

$$\phi_\tau \ W_\tau \ .$$

If the intertemporal elasticity of substitution σ equals 1, which corresponds to log utility, the substitution and income effects exactly cancel. In this case, the interest elasticity of saving depends only on the wealth revaluation effect. If $\sigma = 0$, the substitution effect vanishes (and the rate of time preference δ becomes irrelevant). Note also that the strength of the substitution effect does not depend solely on σ, but also on δ, r, and $(T_2-\tau)$. The parameter σ alone shows how the *slope* of the planned consumption path changes when the interest rate changes, but the *initial jump* to that path--which is the actual change in consumption--depends on these other variables, because they affect the budget constraint that the new consumption path must be consistent with.

The Wealth Effect

The wealth effect includes the revaluation of both human wealth and financial wealth. We begin by calculating the present value of future labor and pension income (human wealth) for an individual of age τ. Suppose that an individual earns M_0 in the first year of his or her working life, and this amount grows over time at rate a until retirement T_1 years later. After that point, the person receives an annuity A (which equals the fraction b of final labor income) each year until T_2 when he or she dies. The present discounted value of labor income can be written as:

$$L_\tau = \lambda_\tau M_\tau = \lambda_\tau e^{a\tau} M_0 \ ,$$

where

$$\lambda_\tau = \frac{e^{(a-r)(T_1-\tau)} - e^{(a-r)[\min(\tau,T_1)-\tau]}}{a-r} \ .$$

The present discounted value of pension income can be written as:

$$P_\tau = \kappa_\tau A = \kappa_\tau b e^{aT_1} M_0 \ ,$$

where

$$\kappa_\tau = \frac{e^{-r[\max(\tau,T_1)-\tau]} - e^{-r(T_2-\tau)}}{r} \ .$$

The revaluation of human wealth when the interest rate changes is given by:

$$\frac{\partial L_\tau}{\partial r} = \alpha_\tau M_\tau = \alpha_\tau e^{a\tau} M_0$$

and

$$\frac{\partial P_\tau}{\partial r} = \beta_\tau A = \beta_\tau b e^{aT_1} M_0 \ ,$$

where

$$\alpha_\tau = \frac{\left(1-(a-r)(T_1-\tau)\right)e^{(a-r)(T_1-\tau)} - \left(1-(a-r)[\min(\tau,T_1)-\tau]\right)e^{(a-r)[\min(\tau,T_1)-\tau]}}{(a-r)^2}$$

and

$$\beta_\tau = \frac{\left(1+r(T_2-\tau)\right)e^{-r(T_2-\tau)} - \left(1+r[\max(\tau,T_1)-\tau]\right)e^{-r[\max(\tau,T_1)-\tau]}}{r^2} .$$

The amount of financial wealth that is held by a person of age τ equals the present discounted value of earnings less the present discounted value of consumption for someone who has been following the lifecycle model (at the same interest rate) for his or her entire life:

$$F_\tau = e^{r\tau}\left[\int_0^{\min(\tau,T_1)} M_t\, e^{-rt}\, dt + \int_{T_1}^{\max(\tau,T_1)} A\, e^{-rt}\, dt - \int_0^{\tau} C_t\, e^{-rt}\, dt\right]$$

$$= e^{r\tau}\left[M_0\,\frac{e^{(a-r)\min(\tau,T_1)}-1}{a-r} + A\,\frac{e^{-rT_1}-e^{-r\max(\tau,T_1)}}{r} - W_0\,\frac{(1-e^{-\theta\tau})\,\mu_0}{\theta}\right].$$

The revaluation of financial wealth when the interest rate changes depends on the form in which the wealth is held. If the wealth is held in bank accounts, then $\partial F_\tau/\partial r$ equals zero. If the wealth is held as shares of physical capital, then:

$$\frac{\partial F_\tau}{\partial r} = -\frac{F_\tau}{r+n} ,$$

where n is the depreciation rate of physical capital.

Appendix B: Numerical Formulas for Target Savers

The sixth section of the paper presents the theory of how target savers respond to a change in the interest rate. This appendix derives numerical formulas for the interest elasticity of saving of the two types of target savers discussed in that section: simple target savers, and people following the advice of financial planners.

Simple Target Saving

The simplest form of target saving involves saving just enough today to accumulate a desired amount by a specific date in the future. This implies that:

$$G_T = S_\tau \, e^{r(T-\tau)} \quad ,$$

where r is the interest rate, τ is the person's current age, T is the person's age when the target must be reached, S_τ is current saving, and G_T is the target amount (or goal). Taking logs and derivatives yields an interest elasticity of saving equal to $-r\,(T-\tau)$. As one would expect, the power of compound interest means that the effect of interest-rate changes is larger (in absolute value) when the target date is further in the future.[70]

A slightly more complicated form of target saving involves saving just enough *for a number of years* to accumulate a desired amount by a specific date in the future. In this case:

[70] Two additional points are worth noting. First, the interest rate used here is an annual rate. If one defines the interest rate instead to be the rate that applies over the entire period until the target date, then the interest elasticity of saving becomes simply $-r$. Second, the elasticity calculated here equals the percent change in saving for a one-percent change in the net rate-of-return r. If one calculated the *semi*-elasticity of saving (the percent change in saving for a one *percentage point* change in r), or the elasticity with respect to the gross return $1+r$, then the elasticity would be $-(T-\tau)$. Making both of these changes together yields a value of -1. This characterization of target saving is the most familiar one to economists.

$$G_{T_4} = e^{r(T_4-\tau)} \int_{\tau}^{T_3} S e^{-r(t-\tau)} dt \quad ,$$

where S is saving each year, T_3 is the person's age when saving stops, and T_4 is the person's age when the target must be reached. The elasticity of saving in this case is:

$$1 - \frac{r \left((T_4-\tau) e^{r(T_4-\tau)} - (T_4-T_3) e^{r(T_4-T_3)} \right)}{e^{r(T_4-\tau)} - e^{r(T_4-T_3)}} \quad .$$

The Financial-Planning Model

Another form of target saving is encouraged by financial planners who recommend saving enough when working to achieve a chosen fraction of pre-retirement consumption during retirement. As explained in the sixth section of the paper, this type of target saver faces the same budget constraint as a lifecycle consumer, but determines the slope of his or her lifetime consumption path differently. As a result, the effect of an interest-rate change on consumption and saving in the financial-planning model can be decomposed into an income effect and a wealth effect. There is no substitution effect because the slope of the individual's consumption path is not affected by the interest rate.

Consumption in the financial-planning model can be described by:

$$C_t = \overline{C} \;\; \text{for all } t < T_1 \quad , \text{ and}$$
$$C_t = \rho \overline{C} \;\; \text{for all } t > T_1 \quad ,$$

where C_t is consumption in period t, ρ is the chosen replacement rate in retirement, and T_1 is the person's age at retirement (measured from the beginning of the working life). The individual's intertemporal budget constraint is:

$$\int_{\tau}^{T_2} C_t e^{-r(t-\tau)} dt = \int_{\tau}^{T_2} Y_t e^{-r(t-\tau)} dt + F_\tau ,$$

where Y_t is the sum of labor and pension income in period t, r is the interest rate, F_τ is financial wealth at time τ, and τ is the person's current age and T_2 the person's age at death (both measured from the beginning of the working life). Substituting the consumption formulas into the budget constraint and denoting the present value of all income (the right side of the budget constraint) as W_τ yields the following formula for consumption while working:

$$\overline{C} = \eta_\tau W_\tau = \frac{r}{1 - (1-\rho)e^{-r[\max(\tau,T_1)-\tau]} - \rho e^{-r(T_2-\tau)}} W_\tau ,$$

where η_τ is the marginal propensity to consume out of wealth.[71] Consumption when retired equals ρ times the value when working.

The derivative of $\overline{C}$ with respect to the interest rate is:

$$\frac{\partial \overline{C}}{\partial r} = W_\tau \frac{\partial \eta_\tau}{\partial r} + \eta_\tau \frac{\partial W_\tau}{\partial r} .$$

Further:

$$\frac{\partial C_t}{\partial r} = \frac{\partial \overline{C}}{\partial r} \quad \textit{for all } t<T_1 \text{ , and}$$
$$\frac{\partial C_t}{\partial r} = \rho \frac{\partial \overline{C}}{\partial r} \quad \textit{for all } t>T_1 .$$

The first term of the expression for $\partial\overline{C}/\partial r$ represents the income effect, and the second term represents the wealth effect. The income effect can be rewritten as:

[71] As an individual grows older, changes in η_τ and W_τ are exactly offsetting, keeping $\overline{C}$ constant.

$$W_\tau \; \frac{\xi_\tau \; - \; r\left((1-\rho)\,[\max(\tau,T_1)-\tau]\, e^{-r[\max(\tau,T_1)-\tau]} \; + \; \rho\,(T_2-\tau)\, e^{-r(T_2-\tau)} \right)}{\xi_\tau^2} \quad ,$$

where

$$\xi_\tau \; = \; 1 \; - \; (1-\rho)\, e^{-r[\max(\tau,T_1)-\tau]} \; - \; \rho\, e^{-r(T_2-\tau)} \quad .$$

The wealth effect includes the revaluation of both human wealth and financial wealth, as in the lifecycle model described in Appendix A. The revaluation of human wealth is the same in the two models. The revaluation of financial wealth is slightly different, because an individual's consumption pattern--and thus his or her accumulated financial wealth at a point in time--is different. The amount of financial wealth that is held by a person of age tau who has been following the financial-planning model (at the same interest rate) for his or her entire life is:

$$e^{r\tau} \left[M_0 \, \frac{e^{(a-r)\min(\tau,T_1)} - 1}{a-r} \; + \; A \, \frac{e^{-rT_1} - e^{-r\max(\tau,T_1)}}{r} \; - \; W_0 \, \frac{\eta_0 \, \zeta_\tau}{r} \right] \quad ,$$

where

$$\zeta_\tau \; = \; 1 \; - \; e^{-r\min(\tau,T_1)} \; + \; \rho\, e^{-rT_1} \; - \; \rho\, e^{-r\max(\tau,T_1)} \quad .$$

Appendix C: Interest Receipts and Payments of Households

The seventh section of the paper examines how changes in interest rates affect household interest receipts and payments. This appendix describes the data and regressions that underlie that analysis.

Data on household interest payments and receipts are drawn from the National Income and Product Accounts (NIPAs). In the first panel of Figure 2, interest receipts equal monetary interest received by persons (Table 8.18, line 29) plus imputed interest received by persons (Table 8.18, line 49). Interest payments equal monetary interest paid by persons (Table 8.18, line 15 or Table 2.1, line 28), and cash flow equals disposable personal income (Table 2.1, line 25) less interest payments. For the second panel of the figure, adjusted interest receipts equal monetary interest received, and adjusted interest payments equal monetary interest paid plus imputed mortgage interest (specifically, net interest for owner-occupied housing and the fixed capital of nonprofit institutions; Table 8.19, lines 93, 100 and 105). Cash flow is not adjusted at this point. The tax adjustments in the third panel of Figure 2 apply a tax rate of 28 percent where appropriate. Thus, tax-adjusted interest receipts equal 0.72 times adjusted interest receipts, and tax-adjusted interest payments equal monetary interest payments plus 0.72 times mortgage interest payments. Tax-adjusted cash flow equals 0.72 times cash flow.

The speed with which household interest receipts and payments adjust to changes in market interest rates is estimated using a procedure suggested by Goodman, Luckett and Wilcox (1988). To begin, I construct implicit average yields using the NIPA data discussed above and data on household assets and liabilities from the Flow of Funds Accounts. The average interest rate on receipts equals monetary interest received divided by the sum of households' deposits and holdings of credit market instruments (Table L.100, lines 2 and 7). The average interest rate on mortgage payments equals imputed mortgage interest divided by households' home mortgage liabilities (Table L.100, line 27). Finally, the average interest rate on other payments equals monetary interest paid divided by households' liabilities of credit market instruments less home mortgages (Table L.100, line 26 minus line 27).

Then, I estimate the following equation for each of these implicit interest rates using annual data from 1962 through 1994:

$$r_t = \omega + \sum_{i=0}^{3} \rho_i rtb_{t-i} + \varepsilon_t ,$$

where r_t is the interest rate in question and rtb_{t-i} is the yield on a one-year Treasury bill. The estimated coefficients and R^2 statistics are shown below, with "*" indicating a coefficient that is statistically different from zero.

Dependent Variable	Independent Variables					R^2
Implicit Interest Rate on:	Constant	Contemporaneous and 3 lags of the Treasury bill rate				
Interest receipts	0.01*	0.27*	0.26*	0.07	0.28*	0.96
Mortgage interest payments	0.03*	0.08	0.23	0.00	0.46*	0.90
Other interest payments	0.08*	0.09	0.07	0.09	-0.01	0.74

References

Orazio P. Attanasio and Martin Browning, "Consumption over the Life Cycle and over the Business Cycle," *American Economic Review*., vol. 85, no. 5 (December 1995), pp. 1118-1137.

Orazio P. Attanasio and Guglielmo Weber, "Consumption Growth, the Interest Rate, and Aggregation," *Review of Economic Studies*, vol. 60, no. 3 (July 1993), pp. 631-649.

Orazio P. Attanasio and Guglielmo Weber, "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, vol. 103, no. 6 (December 1995), pp. 1121-1157.

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

Alan J. Auerbach, Laurence J. Kotlikoff, and David N. Weil, "The Increasing Annuitization of the Elderly--Estimates and Implications for Intergenerational Transfers, Inequality, and National Saving," Working Paper No. 4182 (National Bureau of Economic Research, Cambridge, Mass., October 1992).

Robert B. Avery and Arthur B. Kennickell, "Household Saving in the U.S.," *Review of Income and Wealth*, vol. 37, no. 4 (December 1991), pp. 409-432.

Robert J. Barro, "Are Government Bonds Net Wealth?," *Journal of Political Economy*, vol. 82, no. 6 (December 1974), pp. 1095-1117.

Robert J. Barro, "The Ricardian Approach to Budget Deficits," *Journal of Economic Perspectives*, vol. 3, no. 2 (Spring 1989), pp. 37-54.

Robert B. Barsky, F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Survey," Working Paper No. 5213 (National Bureau of Economic Research, Cambridge, Mass., August 1995).

B. Douglas Bernheim, "A Neoclassical Perspective on Budget Deficits," *Journal of Economic Perspectives*, vol. 3, no. 2 (Spring 1989), pp. 55-72.

B. Douglas Bernheim, "Rethinking Saving Incentives" (Stanford University, January 1996).

B. Douglas Bernheim, Andrei Shleifer, and Lawrence H. Summers, "The Strategic Bequest Motive," *Journal of Political Economy*, vol. 93, no. 6 (December 1985), pp. 1045-1076.

B. Douglas Bernheim and John Karl Scholz, "Private Saving and Public Policy," *Tax Policy and the Economy*, vol. 7 (Cambridge, Mass.: MIT Press, 1993).

Olivier Jean Blanchard and Stanley Fischer, *Lectures on Macroeconomics* (Cambridge, Mass.: MIT Press, 1989).

Michael J. Boskin, "Taxation, Saving and the Rate of Interest," *Journal of Political Economy*, vol. 86, no. 2 (April 1978), pp. S3-S27.

Barry Bosworth, Gary Burtless, and John Sabelhaus, "The Decline in Saving: Evidence from Household Surveys," *Brookings Papers on Economic Activity* (1:1991), pp. 183-241.

Martin Browning and Annamaria Lusardi, "Household Saving: Micro Theories and Micro Facts," (McMaster University, Hamilton, Ontario and Dartmouth College, Hanover, N.H., November 1994).

Ricardo J. Caballero, "Consumption Puzzles and Precautionary Savings," *Journal of Monetary Economics*, vol. 25, no. 1 (January 1990), pp. 113-136.

John Y. Campbell and N. Gregory Mankiw, "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence," in Olivier Jean Blanchard and Stanley Fischer, eds., *NBER Macroeconomics Annual: 1989* (Cambridge, Mass.: MIT Press, 1989), pp. 185-216.

Richard Cantor, "Interest Rates, Household Cash Flow, and Consumer Expenditures," *Quarterly Review*, Federal Reserve Bank of New York, Summer 1989, pp. 59-67.

Christopher D. Carroll, "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity* (2:1992), pp. 61-135.

Christopher D. Carroll and Lawrence H. Summers, "Consumption Growth Parallels Income Growth: Some New Evidence," in B. Douglas Bernheim and John B. Shoven, eds., *National Saving and Economic Performance* (Chicago: The University of Chicago Press, 1991), pp. 305-343.

Congressional Budget Office, *Assessing the Decline in the National Saving Rate*, April 1993.

Angus Deaton, *Understanding Consumption* (Oxford, England: Clarendon Press, 1992).

Robert J. Doyle, Jr. and Eric T. Johnson, *Readings in Wealth Accumulation Planning* (Bryn Mawr, Pennsylvania: The American College, 1991).

Karen E. Dynan, "How Prudent Are Consumers?" *Journal of Political Economy*, vol. 101, no. 6 (December 1993), pp. 1104-1113.

Karen E. Dynan, "Relative Wage Changes and Estimates of the Rate of Time Preference" (Federal Reserve Board, September 1994).

Eric M. Engen, "Consumption and Saving in a Life-Cycle Model with Stochastic Earnings and Mortality Risk" (Federal Reserve Board, March 1992).

Eric M. Engen, William G. Gale, and John Karl Scholz, "Do Saving Incentives Work?" *Brookings Papers on Economic Activity* (1:1994), pp. 85-151.

Owen J. Evans, "Tax Policy, the Interest Elasticity of Saving, and Capital Accumulation: Numerical Analysis of Theoretical Models," *American Economic Review*, vol. 73, no. 3 (June 1983), pp. 398-410.

Martin Feldstein, "Fiscal Policies, Capital Formation, and Capitalism," Working Paper 4885 (National Bureau of Economic Research, Cambridge, Mass., October 1994).

Irwin Friend and Joel Hasbrouck, "Saving and After-Tax Rates of Return," *The Review of Economics and Statistics*," vol. 65, no. 4 (November 1983), pp. 537-543.

William G. Gale and John Karl Scholz, "IRAs and Household Saving," *American Economic Review*, vol. 84, no. 5 (December 1994), pp. 1233-1260.

William G. Gale and John Karl Scholz, "Intergenerational Transfers and the Accumulation of Wealth," *Journal of Economic Perspectives*, vol. 8, no. 4 (Fall 1994), pp. 145-160.

John L. Goodman, Jr., Charles A. Luckett, and David W. Wilcox, "Interest Rates and Household Cash Flow" (Federal Reserve Board, December 1988).

Jane G. Gravelle, *The Economic Effects of Taxing Capital Income* (Cambridge, Mass.: MIT Press, 1994).

Robert E. Hall, "Stochastic Implications of the Life Cycle - Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, vol. 86, no. 6 (December 1978), pp. 971-987.

Robert E. Hall, "Intertemporal Substitution in Consumption," *Journal of Political Economy*, vol. 96, no. 2 (April 1988), pp. 339-357.

Patric H. Hendershott, "Comment," in Joel Slemrod, ed., *Do Taxes Matter?: The Impact of the Tax Reform Act of 1986* (Cambridge, Mass.: MIT Press, 1990).

Philip F. Howrey and Saul H. Hymans, "The Measurement and Determination of Loanable Funds Savings," *Brookings Papers on Economic Activity*, vol. 3 (1978), pp. 655-685.

R. Glenn Hubbard, Jonathan Skinner, and Stephen P. Zeldes, "Precautionary Saving and Social Insurance," *Journal of Political Economy*, vol. 103, no. 2 (April 1995), pp. 360-399.

Larry E. Jones, Rodolfo E. Manuelli, and Peter E. Rossi, "Optimal Taxation in Models of Endogenous Growth, *Journal of Political Economy*, vol. 101 (June 1993), pp. 485-517.

Robert G. King and Sergio Rebelo, "Public Policy and Economic Growth: Developing Neoclassical Implications," *Journal of Political Economy*, vol. 98, no. 5 (October 1990), pp. S126-S150.

Laurence J. Kotlikoff, "Intergenerational Transfers and Savings," *Journal of Economic Perspectives*, vol. 2, no. 2 (Spring 1988), pp. 41-58.

Laurence J. Kotlikoff and Lawrence H. Summers, "The Role of Intergenerational Transfers in Aggregate Capital Accumulation," *Journal of Political Economy*, vol. 89, no. 4 (August 1981), pp. 706-732.

Emily C. Lawrance, "Poverty and the Rate of Time Preference: Evidence from Panel Data," *Journal of Political Economy*, vol. 99, no. 1 (February 1991), pp. 54-77.

George Loewenstein and Drazen Prelec, "Negative Time Preference," *American Economic Review*, vol. 81, no. 2 (May 1991), pp. 347-352.

Robert E. Lucas, "Supply-Side Economics: An Analytical Review," *Oxford Economic Papers*, vol. 42 (April 1990), pp. 293-316.

N. Gregory Mankiw, "Consumer Spending and the After-Tax Real Interest Rate," in Martin Feldstein, ed., *The Effects of Taxation on Capital Accumulation* (Chicago: The University of Chicago Press, 1987), pp. 53-68.

Joyce M. Manchester and James M. Poterba, "Second Mortgages and Household Saving," *Regional Science and Urban Economics*, vol. 19, no. 2 (1989), pp. 325-346.

Eileen Mauskopf, "The Transmission Channels of Monetary Policy: How Have They Changed?," *Federal Reserve Bulletin*, vol. 76, no. 12 (December 1990), pp. 985-1008.

Franco Modigliani, "The Role of Intergenerational Transfers and Life Cycle Saving in the Accumulation of Wealth," *Journal of Economic Perspectives*, vol. 2, no. 2 (Spring 1988), pp. 15-40.

Mancur Olson and Martin J. Bailey, "Positive Time Preference," *Journal of Political Economy*, vol. 89, no. 1 (February 1981), pp. 1-25.

James M. Poterba, "Are Consumers Forward-Looking? Evidence from Fiscal Experiments," *American Economic Review*, vol. 78, no. 2 (May 1988), pp. 413-418.

James M. Poterba, "Personal Saving Behavior and Retirement Income Modelling: A Research Assessment" (Massachusetts Institute of Technology, November 1994).

James M. Poterba, Steven F. Venti, and David A. Wise, "Do 401(k) Contributions Crowd Out Other Personal Saving?," *Journal of Public Economics*, vol. 58, no. 1 (August 1995), pp. 1-32.

William C. Randolph and Diane Lim Rogers, "The Implications for Tax Policy of Uncertainty About Labor Supply and Savings Responses," *National Tax Journal*, vol. 48, no. 3 (September 1995), pp. 429-446.

David E. Runkle, "Liquidity Constraints and the Permanent-Income Hypothesis: Evidence from Panel Data," *Journal of Monetary Economics*, vol. 27, no. 1 (February 1991), pp. 73-98.

Matthew D. Shapiro, "The Permanent Income Hypotheses and the Real Interest Rate: Some Evidence from Panel Data," *Economics Letters*, vol. 14 (1984), pp. 93-100.

Hersh M. Shefrin and Richard H. Thaler, "The Behavioral Life-Cycle Hypothesis," *Economic Inquiry*, vol. 26, no. 4 (October 1988), pp. 609-643.

Jonathan Skinner and Daniel Feenberg, "The Impact of the 1986 Tax Reform on Personal Saving," in Joel Slemrod, ed., *Do Taxes Matter?: The Impact of the Tax Reform Act of 1986* (Cambridge, Mass.: MIT Press, 1990).

Lawrence H. Summers, "Capital Taxation and Accumulation in a Life Cycle Growth Model," *American Economic Review*, vol. 71, no. 4 (September 1981), pp. 533-544.

Lawrence H. Summers, "Tax Policy, the Rate of Return, and Savings," Working Paper No. 995 (National Bureau of Economic Research, Cambridge, Mass., 1982).

Lawrence H. Summers, "The After-Tax Rate of Return Affects Private Savings," *American Economic Review*, vol. 84, no. 2 (May 1984), pp. 249-253.

Richard H. Thaler, "Psychology and Savings Policies," *American Economic Review*, vol. 84, no. 2 (May 1994), pp. 186-192.

Steven F. Venti and David A. Wise, "Have IRAs Increased U.S. Savings?: Evidence from Consumer Expenditure Surveys," *Quarterly Journal of Economics*, vol. 105, no. 3 (August 1990), pp. 661-698.

David W. Wilcox, "Social Security Benefits, Consumption Expenditure, and the Life Cycle Hypothesis," *Journal of Political Economy*, vol. 97, no. 2 (April 1989), pp. 288-304.

David W. Wilcox, "Interest Rates and Consumption: A Classical Analysis" (Federal Reserve Board, March 1993).

Stephen P. Zeldes, "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, vol. 97, no. 2 (April 1989), pp. 305-346.

Table 1
Estimates of the Intertemporal Elasticity of Substitution

Study	Estimate (standard error)	Consumption data; Rate-of-return data[a]
Hall (1988)	0.35 (.34) 0.27 (.33) 0.07 (.05)	semiannual NIPA data; Treasury bills, savings accounts and stocks, respectively
	-0.40 (.20)[b]	annual NIPA data; Treasury bills
	-0.03 (.38)[b] 0.03 (.10)	monthly NIPA data; Treasury bills and stocks, respectively
	0.10 (.23)	quarterly NIPA data; Treasury bills
Campbell and Mankiw (1989)	0.14[c] 0.17[c] 0.05[c]	quarterly NIPA data; Treasury bills[d]
Shapiro (1984)	26.5 (28.6) 20.4 (27.1)	PSID data for all households; Treasury bills[e]
Zeldes (1989)	0.37 (.24) -1.46 (.36) 1.92 (.81)	PSID data for households with low wealth-income ratios; Treasury bills[f]
	0.43 (.31) 1.44 (1.34) 0.58 (.32)	PSID data for households with high wealth-income ratios; same[f]
Runkle (1991)	0.45 (.16) 0.35 (.16) 0.49 (.20) 0.49 (.17) 0.49 (.17) 0.43 (.16) 0.48 (.17) 0.48 (.17)	PSID data for all households; savings accounts[d]

	0.30 (.23) 0.28 (.22) 0.23 (.22) 0.28 (.22) 0.31 (.22)	PSID data for households with low wealth-income ratios; same[d]
	0.68 (.27) 0.67 (.26) 0.54 (.25) 0.67 (.26) 0.56 (.25)	PSID data for households with high wealth-income ratios; same[d]
Lawrance (1991)	0.54 (.40) 0.73 (.70)	PSID data for all households; Treasury bills and savings accounts, respectively[g]
	0.54 (.88) 0.46 (1.50)	PSID data for households with below-median income; same[g]
	0.81 (.49) 1.80 (.89)	PSID data for households with above-median income; same[g]
Dynan (1993)	0.10 (.09) 0.11 (.85) 0.12 (.87)	CEX data for all households; Treasury bills
Attanasio and Weber (1993)	0.35 (.13) 0.33 (.10)	British National Accounts data; savings accounts[e]
	0.37 (.18) 0.29 (.15) 0.38 (.21) 0.29 (.25)	FES data for all households; same[e]
	0.60 (.33) 0.63 (.28) 0.78 (.28) 0.73 (.27)	FES data for middle-aged cohort; same[e]
	0.56 (.20)	FES data for young cohort; same
	0.32 (.36)	FES data for old cohort; same

Attanasio and Weber (1995)	0.39 (.28) 0.34 (.28) 0.15 (.35) 0.39 (.21) 0.48 (.28) 0.33 (.32) 0.56 (.22) 0.67 (.19)	CEX data for all households; municipal bonds[e]

Notes:

[a] The studies use the real after-tax rate-of-return. PSID is the Panel Study of Income Dynamics, CEX is the Consumer Expenditure Survey, and FES is the British Family Expenditure Survey. Only estimates based on instrumental variables are included because of the endogeneity that arises in estimating Euler equations.

[b] Hall notes that a negative estimate "cannot be taken literally since it implies nonconcave utility." Instead, he concludes that "the case for a significantly positive value ... cannot be made" here (p. 353).

[c] In this model, the elasticity is a nonlinear function of the coefficient on the interest rate and the coefficient on contemporaneous income. Standard errors for the elasticity estimates are not available, but they are presumably fairly large because the coefficient on the interest rate alone is always less than its standard error.

[d] Different estimates arise from alternative instruments.

[e] Different estimates arise from alternative specifications.

[f] Different estimates arise from alternative ways of splitting the sample.

[g] Only estimates that include time dummies to control for aggregate shocks to consumption are included.

Table 2
Short-Run Interest Elasticity of Saving in the Lifecycle Model

intertemporal elasticity of substitution = 0
rate of time preference = 0
annual rate of wage growth = 0

		Age		
		30 years	50 years	70 years
Substitution effect		0	0	0
Income effect		-7.28	-2.30	-0.41
Human wealth effect		6.13	1.37	0.13
Financial wealth effect:	bank accounts	0	0	0
	physical capital	0.08	0.33	0.54
Total elasticity:	bank accounts	-1.16	-0.93	-0.28
	physical capital	-1.08	-0.60	0.26

Table 3
Short-Run Interest Elasticity of Saving in the Lifecycle Model

intertemporal elasticity of substitution = 0.33
rate of time preference = 0
annual rate of wage growth = 0

		Age		
		30 years	50 years	70 years
Substitution effect		0.68	0.57	0.13
Income effect		-2.06	-1.73	-0.39
Human wealth effect		1.45	0.76	0.08
Financial wealth effect:	bank accounts	0	0	0
	physical capital	0.07	0.44	0.58
Total elasticity:	bank accounts	0.07	-0.41	-0.18
	physical capital	0.13	0.03	0.40

Table 4
Short-Run Interest Elasticity of Saving in the Lifecycle Model

intertemporal elasticity of substitution = 0.33
rate of time preference = -0.02
annual rate of wage growth = 0

		Age		
		30 years	50 years	70 years
Substitution effect		0.48	0.49	0.13
Income effect		-1.46	-1.48	-0.38
Human wealth effect		0.94	0.56	0.07
Financial wealth effect:	bank accounts	0	0	0
	physical capital	0.06	0.43	0.59
Total elasticity:	bank accounts	-0.04	-0.43	-0.19
	physical capital	0.02	-0.01	0.40

Table 5
Short-Run Interest Elasticity of Saving in the Lifecycle Model

intertemporal elasticity of substitution = 0.33
rate of time preference = 0.02
annual rate of wage growth = 0

		Age		
		30 years	50 years	70 years
Substitution effect		1.08	0.67	0.13
Income effect		-3.28	-2.02	-0.40
Human wealth effect		2.52	1.02	0.10
Financial wealth effect:	bank accounts	0	0	0
	physical capital	0.07	0.41	0.56
Total elasticity:	bank accounts	0.32	-0.33	-0.16
	physical capital	0.39	0.09	0.40

Table 6
Short-Run Interest Elasticity of Saving in the Lifecycle Model

intertemporal elasticity of substitution = 0.33
rate of time preference = 0
annual rate of wage growth = 0.01

		Age		
		30 years	50 years	70 years
Substitution effect		1.28	0.58	0.14
Income effect		-3.88	-1.77	-0.42
Human wealth effect		3.05	0.94	0.11
Financial wealth effect:	bank accounts	0	0	0
	physical capital	0.06	0.30	0.58
Total elasticity:	bank accounts	0.45	-0.24	-0.17
	physical capital	0.51	0.06	0.41

Table 7
Short-Run Interest Elasticity of Saving in the Lifecycle Model

intertemporal elasticity of substitution = 0.33
rate of time preference = 0
annual rate of wage growth = 0.03

		Age		
		30 years	50 years	70 years
Substitution effect		0	0.83	0.17
Income effect		0	-2.51	-0.52
Human wealth effect		0	1.81	0.20
Financial wealth effect:	bank accounts	0	0	0
	physical capital	0	0.07	0.60
Total elasticity:	bank accounts	0	0.13	-0.15
	physical capital	0	0.21	0.45

Table 8
Aggregate Short-Run Interest Elasticity of Saving in the Lifecycle Model

intertemporal elasticity of substitution = 0.33
rate of time preference = 0
annual rate of wage growth = 0.03

	Combined Productivity and Population Growth Rates Across Annual Cohorts		
	2 percent	3 percent	4 percent
Initial Saving Rate	0.05	0.06	0.07
All age groups hold financial wealth in the form of physical capital			
Change in Saving Rate Induced by a 1 Percentage Point Rise in the Interest Rate	0.02	0.02	0.01
Short-Run Interest Elasticity of Saving	1.42	0.99	0.77
Older groups hold a smaller share of wealth as physical capital and a larger share as bank accounts			
Change in Saving Rate Induced by a 1 Percentage Point Rise in the Interest Rate	0.01	0.01	0.01
Short-Run Interest Elasticity of Saving	0.89	0.65	0.54

Table 9
Short-Run Interest Elasticity of Saving in the Financial-Planning Model

replacement rate = 0.8
annual rate of wage growth = 0

		Age		
		30 years	50 years	70 years
Substitution effect		0	0	0
Income effect		-9.81	-3.02	-0.46
Human wealth effect		8.62	2.02	0.18
Financial wealth effect:	bank accounts	0	0	0
	physical capital	0.08	0.35	0.54
Total elasticity:	bank accounts	-1.19	-1.00	-0.28
	physical capital	-1.11	-0.65	0.26

Table 10
Aggregate Short-Run Interest Elasticity of Saving in the Financial-Planning Model

replacement rate = 0.8
annual rate of wage growth = 0

	Combined Productivity and Population Growth Rates Across Annual Cohorts		
	2 percent	3 percent	4 percent
Initial Saving Rate	0.04	0.05	0.06
All age groups hold financial wealth in the form of physical capital			
Change in Saving Rate Induced by a 1 Percentage Point Rise in the Interest Rate	-0.01	-0.01	-0.01
Short-Run Interest Elasticity of Saving	-1.00	-0.90	-0.87
Older groups hold a smaller share of wealth as physical capital and a larger share as bank accounts			
Change in Saving Rate Induced by a 1 Percentage Point Rise in the Interest Rate	-0.02	-0.02	-0.02
Short-Run Interest Elasticity of Saving	-1.48	-1.17	-1.06

Figure 1
Effects on the Budget Constraint of an Increase in the Interest Rate

The substitution effect results from the pivot of the budget constraint. The income effect results from the change in the present discounted value of consumption labelled "A." The human wealth effect results from the change in the present discounted value of labor income labelled "B."

Case A: All labor income is earned in the first period of life, so the budget constraint is: $(1+r)\ C_1 + C_2 = (1+r)\ Y_1$

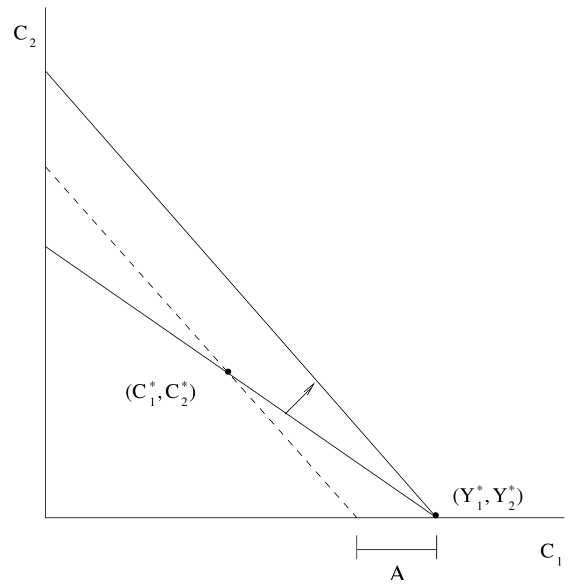

Case B: Labor income is earned in both the first and second periods of life, so the budget constraint is: $(1+r)\ C_1 + C_2 = (1+r)\ Y_1 + Y_2$

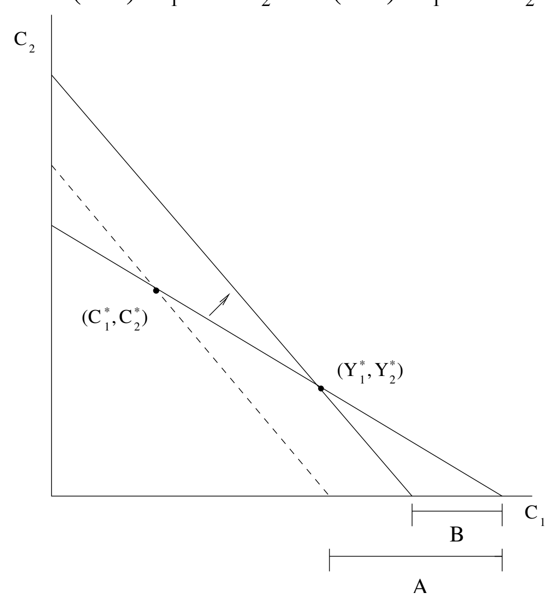

Figure 2
Interest Receipts and Payments of Households

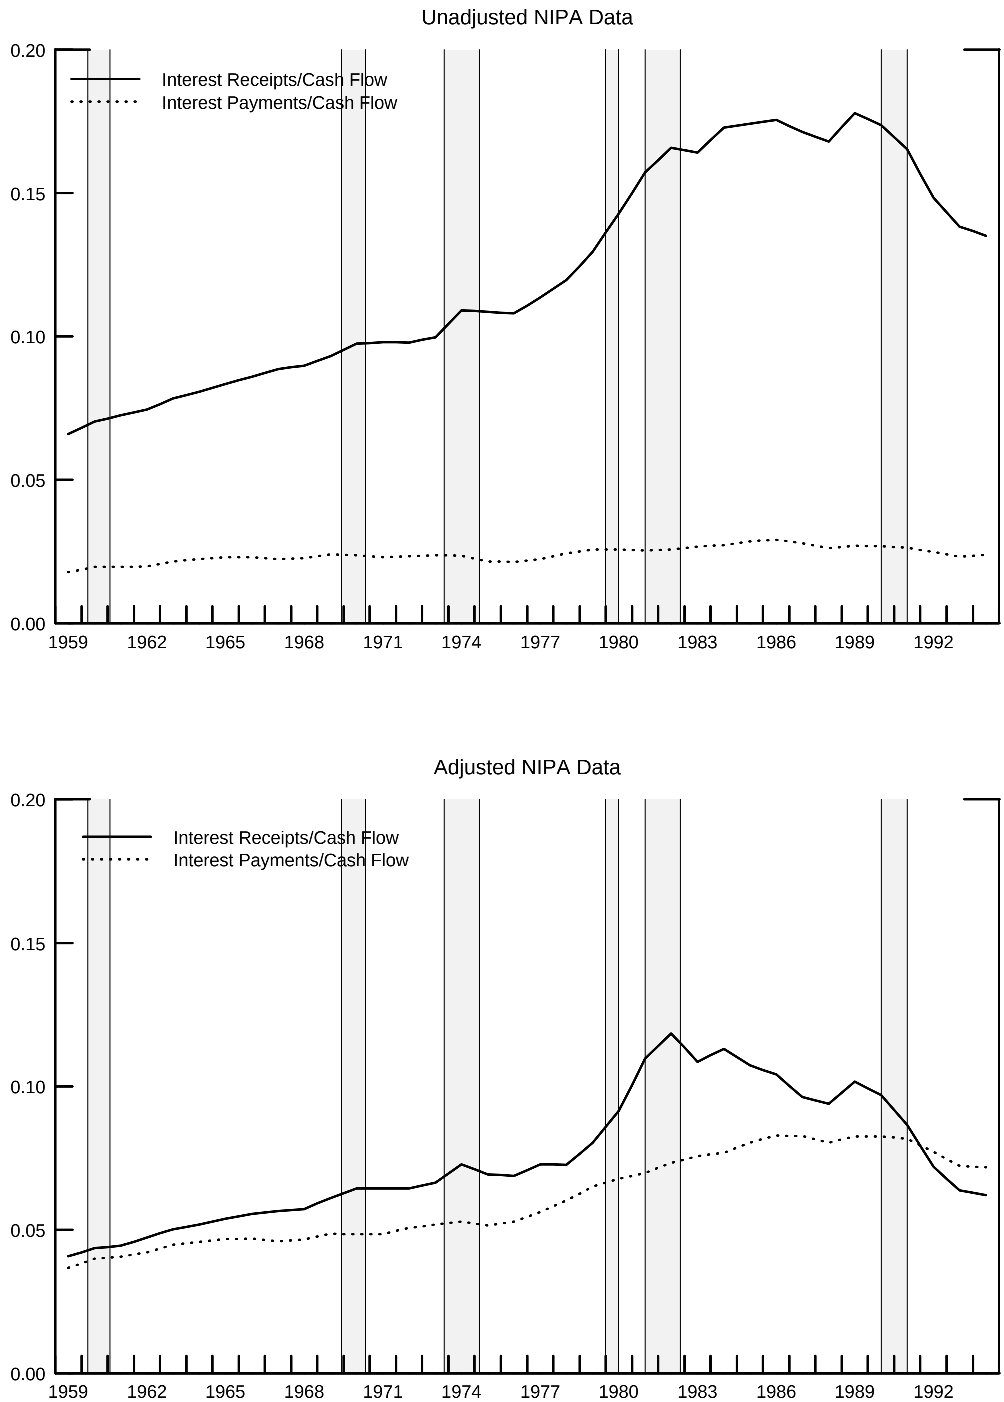

Figure 2 (continued)
Interest Receipts and Payments of Households