NBER WORKING PAPER SERIES

SPENDTHRIFT IN AMERICA?
ON TWO DECADES OF DECLINE
IN THE U.S. SAVING RATE

Jonathan A. Parker

Working Paper 7238
http://www.nber.org/papers/w7238

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 1999

For useful comments and discussions, I am grateful to Ben Bernanke, Angus Deaton, Michael Hovarth, Rodolfo Manuelli, Julio Rotemberg, John Karl Scholz, participants at the NBER Macroeconomics Annual Conference, particularly my discussants, and especially Pierre Olivier Gourinchas and Kenneth West. Eric Hurst and Joe Lupton provided invaluable consultations on using the early release PSID data. I thank Karen Dynan for sharing her understanding of NIPA saving measures. Grigori Kosenok provided excellent research assistance with the aggregate series. I alone am responsible for any errors. All opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

Spendthrift in America? On Two Decades of
Decline in the U.S. Saving Rate
Jonathan A. Parker
NBER Working Paper No. 7238
July 1999
JEL No. E2, D1, D9

ABSTRACT

During the past two decades, the personal saving rate in the United States has fallen from eight percent to below zero. This paper demonstrates that this change represents a major shift in the allocation of newly produced goods. The share of GDP that households consume rose by 6 percentage points since 1980. This increase occurred concurrently with a reduction in the growth rate of real consumption spending per person, high real rates of return, and an increasing ratio of aggregate wealth to income. Despite this last fact, wealth changes can explain little of the boom in consumption spending. The largest increases in national wealth post-date the consumption boom and households with different wealth levels have similar increases in consumption. The paper also finds that the changing age distribution of the U.S. population does not explain the consumption boom. While it may be that new wealthier cohorts are driving this boom, the preponderance of evidence suggest rather that the rising consumption to income ratio is due to a common time effect. The main findings of the paper are consistent with either an increase in the discount rate or with a general belief in better economic times in the future. Alternatively, the low rates of saving could be due to a combination of factors such as the increase in intergenerational transfers from the Social Security system raising the consumption of the elderly and an increase in access to credit and expanded financial instruments raising the consumption of the young.

Jonathan A. Parker
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706
jparker@ssc.wisc.edu

Contents

1. Introduction

During the past two decades, the personal saving rate in the United States has fallen dramatically. From a typical and quite steady level of around eight percent during the sixties and seventies, the personal savings rate has declined to below two percent in 1997 and preliminary estimates put the rate at one-half percent in 1998 and negative so far in 1999. Figure 1*a* displays the U.S. personal saving rate from 1959 to 1998 and makes clear the magnitude of the change.[1]

This change does not merely reflect labelling or measurement issues. In particular, for the majority of this decline, it is not the case that businesses or governments have increased their saving with national saving unaffected. Since the National Income and Product Account ($NIPA$) definitions of savings rates are neither transparent nor representative of basic economic concepts, the simplest way to judge the importance of this shift in the U.S. economy is to examine whether consumption has risen as a share of national output. Figure 1*b* shows that the decline in personal saving has largely been mirrored on the expenditure side of the national accounts. The ratio of consumption to GDP in the United States was roughly constant from 1950 to 1980, and has risen by 6 percentage points during the past two decades.[2]

While the ratio of consumption to income has risen significantly, it is worth noting that this ratio has not risen in the past few years. The recent decline in the personal saving rate that has received so much attention from journalists and policymakers is not reflected in the ratio of consumption to output. As the next section shows, the decline in the personal saving rate over the past five years is more than offset by increases in saving by governments and businesses.

This paper focuses on the fundamental and significant change in the allocation of the output of the US economy documented in figure 1*b*: why has the largest economy in the world over twenty years increased its consumption expenditures by six percent of output? This change poses a basic challenge to economists as those who seek to explain economic outcomes. Do we understand the allocation of resources?

This consumption boom also has import for the economic future of the United States. Saving is the accumulation of resources on which to base future consumption. Absent offsetting changes in the national economy, higher consumption generally leads to a lower capital stock and thus adversely impacts wages and national output in the future. If the present low saving rate represents an optimal response of well-functioning markets to fundamental improvements such as new technolo-

[1] The personal saving rate is defined as one minus the ratio of personal outlays to disposable income. In the national accounts, personal outlays are personal consumption expenditures plus interest paid by persons and personal transfer payments abroad; disposable income is labor income, proprietors' income, rental income, personal interest and dividend income, and transfer payments to persons all less personal contributions for social insurance and personal tax and non-tax payments.

[2] The economic significance of all changes presented throughout the paper are not figments of idiosyncratic start or end dates.

gies, then policies designed to stimulate saving are at best unnecessary, since the future is rosy. On the other hand, if high consumption rates are the results of imprudent fiscal policies or malfunctioning markets, then anemic saving signals an avoidably worse future.[3]

Despite the basic prediction that lower saving and investment lowers the capital stock, the aggregate wealth to income ratio has actually increased during the period of the consumption boom. While new investment has slowed, the revaluation of existing assets has kept wealth levels high relative to national output, raising the possibility that the capital to income ratio is not declining. Put differently, the saving rate including capital gains has not fallen. As is clear from the paper and comments on this topic in this Volume, however, high stock prices may not signal high future dividends or a valuable capital stock. It seems imprudent to simply assume that saving is in some sense high and that the capital stock is larger than would be inferred from past investment. Instead, this paper considers whether the appreciation of assets, whatever the driving mechanism, can explain the changing allocation of current output. If it can, this suggests that households perceive the increase in the value of the stock market as real wealth creation. As this paper demonstrates however, the increase in wealth alone does not explain the consumption boom.

This paper begins by laying out the basic facts surrounding the decline in national saving and how a canonical aggregate model can account for these changes. I focus both on several recently observed changes in the U.S. economy and on the main extant theories of the increased consumption of output.[4]

I employ two main sources of data to study the increase in the consumption of output. First, the paper uses U.S. national accounts data to compare the timing of the consumption boom with the timing of the candidate driving forces, and to ask what expected changes would be required to rationalize observed household consumption behavior. Second, the work evaluates cross-sectional implications of the theories using a custom-built panel dataset on U.S. households. I impute consumption of nondurable goods and services for each household in the Panel Study of Income Dynamics ($PSID$) using information from the Consumer Expenditure Survey (CEX) and the U.S. national accounts. The resulting data contains $80,000$ observations on household income, consumption, wealth, and demographic characteristics covering the period of interest.

The analysis leads to the following main results.[5]

[3] While uncovering significant evidence about the behavior of the consumption ratio, this paper does not enter this debate directly. For examples these arguments see Bernheim and Shoven, eds (1991), Bernheim and Scholz (1993), Gustman and Steinmeier (1998).

[4] There are no shortage of theories that can "explain" the decline in saving, once one allows any combination of changing structural parameters or shifting definitions as plausible candidates. This paper limits the scope of its investigation to the main extant theories and looks at the data with these first-pass explanations in mind. Further, this paper focuses on ruling out monocausal explanations and upon describing behavior.

[5] While much relevant literature is cited where appropriate, the literature is too large to cover in detail here. See Browning and Lusardi (1996), Hayashi (1997), and Attanasio (1997b) for an

First, the decline in measured saving is not purely due to a rise in expenditures without an associated rise in consumption. That is, households are not simply spending more on durable goods and thereby shifting the composition of their savings.

Second, the consumption boom cannot be explained by decreased government purchases "crowding in" consumption. The sum of government and household expenditures on goods and services has also risen over this period. Further, the declines in government spending that would have to be expected to rationalize the consumption boom are, to the author, implausibly large.

Third, the data suggest that at most one-fifth of the increase in consumption to income can be explained by changes in the ratio of household wealth to income. The consumption boom precedes the recently observed increases in wealth, and the national saving rate has actually risen coincident with the stock market boom of the late 1990's. Additionally, the increases in consumption to income ratios across groups are not related to the distribution of wealth, home ownership, or pension participation. Shocks to asset values are not the main force driving the relative increase in consumption.

Fourth, during this period of rising consumption share, the growth rate of real consumption per capita was low and real interest rates were relatively high. Absent a run of expectational errors, the consumption Euler equation implies that the actual or effective discount rate of the representative agent was high. Additionally, there is a strong correlation between the real interest rate and consumption growth within the period of consumption boom. That is, the aggregate consumption Euler equation provides a better description of the data during this period than in previous periods.

Fifth, turning to evaluating explanations that are consistent with such increased impatience, the changing age distribution and income-by-age distribution of the population are not important causes of the consumption boom. Nor, sixth, can financial innovation which relaxes liquidity constraints and potentially reduces precautionary saving be blamed for the consumption boom. Given the observed increases in debt, this source can generate only one third of the increase in consumption observed to date.

Seventh, the consumption to income ratio of each generation is larger than the generation before them.[6] This implies that intergenerational fiscal transfers alone cannot account for the decline in saving. Thus either different factors have increased the consumption of different generations or general optimism or a preference shift has increased the consumption to income ratios of all households.[7]

In sum, the analysis reveals that each of the major extant theories of the decline

overview of the state of empirical research on saving.

[6] As will be shown, this can be explained either by a time effect increasing everyone's consumption to income or by true cohort effects, as is described here. In this latter case, while the changing age distribution of the population is irrelevant, who is at each age is very relevant.

[7] An example of such a combination of factors is federal transfers from future generations to the elderly and financial innovation that allow the young to consume more out of future income.

in the U.S. saving rate fails on its own to match significant aspects of the macroeconomic or household data. The concluding section of the paper presents some hypotheses that are consistent with the stylized facts uncovered in this paper and with the limited roles found for the monocausal explanations.

The paper is organized as follows. The next section describes the history of the decline in the personal saving rate and its relationship to the allocation of output. Section 3 presents a canonical aggregate model and the classes of explanations for the consumption boom that the paper considers. Section 4 evaluates a subset of the theories using *U.S.* national accounts data and in doing so provides a more detailed description of the aggregate facts. Section 5 describes the main features and construction of the household-level dataset that is used to further test the theories in sections 6, 7, and 8. These sections differ by methodology: section 6 decomposes the consumption to income ratios into age, time and cohort effects; section 7 models the cohort effect and estimates consumption functions; section 8 estimates Euler equations. Section 9 concludes. A data appendix is provided.

2. The Decline in the U.S. Saving Rates

Before turning to the theoretical determinants of the consumption ratio and evaluating these determinants using the aggregate data, this section presents the stylized facts concerning the declining U.S. savings rates.[8] It is important to clarify what has occurred before turning to possible explanations. The section is structured as about saving since it is national saving (and international capital flows) that equals total national investment.

Is the precipitous decline in personal saving shown Figure 1*a* leading to lower national saving, or is public and business saving offsetting the decline? Actually, from private saving data, it is not even clear that households themselves are saving less. National Account data misallocates several categories of saving between private and business saving. Personal saving includes the saving of non-corporate, non-financial businesses such as sole proprietorships, partnerships and nonprofit organizations, that might be better included in business saving. Additionally, because disposable personal income includes nominal rather than real interest payments to businesses, personal saving is overstated relative to business saving.[9]

Given that personal saving is confounded with business saving, the first question is what has happened to their sum, private saving. Figure 2*a* displays the private saving rate — the ratio of private saving to national income — over the last forty years.[10] Prior to the precipitous decline in personal saving, the private saving

[8] In contemporaneous research, Gale and Sabelhaus (forthcoming) analyze the aggregate data on saving and wealth and reach similar conclusions to those of this section.

[9] See Hendershott and Peek (1988) and Summers and Carroll (1987).

[10] NIPA saving rate measures have recently been revised so as to exclude the capital gains distributions of mutual funds from both saving and disposable income. This is consistent with the national accounts purpose to describe the allocation of newly-produced, final, value-added. Unfortunately, this revision only goes back to 1982, so that there is a break in the savings series

rate was close to constant. This stable relationship was known as Denison's Law (Denison (1958)), and this law appears to have been repealed.[11]

Turning now to government saving, higher saving by the government, holding expenditures constant, leads to lower taxes in the future. The principal of Ricardian equivalence states that if taxes are non-distortionary, this offset is complete: households observing higher government saving save less themselves, holding government purchases of goods constant. Figure 2*a* also shows that government saving– the difference between private and gross saving– declined through the 1950's, 1960's and 1970's and has only rebounded from near zero in the early 1990's.[12] Thus gross saving declined steadily from the late 1960's to the early 1990's and has risen recently. We can conclude that while the last five years of declining private saving has been offset by increased government saving, national saving has still fallen substantially in the last twenty years.

One reason for pausing to examine national saving — and not simply focussing on consumption to income ratios throughout — is that saving and investment have moved in lock-step over most of the post-War period. Capital inflows have not offset the decline in saving either because of an offsetting temporal pattern of changes in the world economy or because of any one of the proposed rationalizations of the Feldstein-Horiaka puzzle. Nevertheless, declining national saving has been associated with a large decline in new investment as a share of GDP over the past twenty years. Ultimately, and ceteris paribus, one would expect this decline of roughly one fifth in gross investment to lower the U.S. capital stock per worker by one fifth.

As discussed in the introduction, the value of extant assets, however, has not declined. The stock of wealth in the U.S. has risen as a share of income over the last twenty years. Figure 2*b* shows that net worth, as measured in the Flow of Funds data, has increased as a share of disposable income during the period of declining saving. While the flow of the share of output stored for future production is declining over time , the value of the stock is rising. Put slightly differently, while active saving has decreased, the change in household wealth as a share of income has increased.[13]

These coincident trends raise two puzzles. First, why has wealth risen while

in that year. However, these distributions account for only one-quarter of one percentage point of the saving rate in the 1980's. Thus, while this revision lowered measured private saving in the 1990's significantly, carrying the revision back farther would have a trivial impact on measured saving rates and the conclusions of the present analysis.

[11] In part, Denison's Law is also based on an observed high negative correlation between personal and business saving. Hendershott and Peek (1988) argue that mismeasurement generates most if not all of this negative correlation and thus that Denison's law was never passed in the first place.

[12] This is the official measure of government saving, which does not include changes in government debts associated with Social Security and the implicit and violable promises to future generations. Officially, government obligations held by the public rose by about twenty percent of GDP from 1979 to 1995 and have declined by about five percent since.

[13] This fact is also present in the PSID data that will be used subsequently (Hurst, Luoh and Stafford 1998).

saving has fallen? This question is addressed elsewhere in this volume. Second, and the focus of the balance of this paper, what has driven the decline in active saving and the increasing consumption of output?

3. The Canonical Theory and Main Explanations

This section discusses the main explanations for the consumption boom in the context of a canonical macroeconomic model. In subsequent sections, this explanations are evaluated using aggregate time series evidence and then panel data on household behavior.

To provide a framework for what can be learned about the decrease in saving and the increase in the consumption share of output, I begin with a standard Ramsey economy. Aggregate output, Y, is produced from the aggregate capital stock, K, and total labor in the economy, N, using a constant returns-to scale production technology

$$Y = F(K, AN)$$

where A is exogenous Harrod-neutral technology that grows at rate a. Let the labor force grow at exogenous rate n and let capital depreciate at rate δ. Then one can rearrange the standard capital accumulation equation to solve for the consumption share of output:

$$\frac{c}{y} = 1 - g - (n + a + \delta)\frac{k}{f(k)} - \frac{\dot{k}}{f(k)} \tag{3.1}$$

where C is aggregate consumption, lower case letters denote per-effective-worker values (e.g. $c \equiv \frac{C}{AN}$), g is the rate of government consumption of output, and $f(k) \equiv F(k, 1)$. In steady state, the consumption ratio is related only to the accumulated capital stock, the share of output consumed by the government, and the exogenous rates of technology growth, population growth, and depreciation.

In the canonical Ramsey model with a single infinitely-lived representative agent maximizing the present discounted value of utility flows, the steady-state real interest rate and thus the capital-output ratio are tied down by the modified golden rule. Assuming a Cobb-Douglas production function, the consumption share of output in steady state is

$$\frac{c}{y} = 1 - g - \frac{\alpha(n + a + \delta)}{r + \delta} \tag{3.2}$$

$$= 1 - g - \frac{\alpha(n + a + \delta)}{\frac{a}{\sigma} + \rho + \delta} \tag{3.3}$$

where r is the real interest rate, ρ is the discount factor of the representative agent, α is the share of output that is paid to capital, and σ is the intertemporal elasticity of substitution of the representative agent. The consumption ratio is increased by increases in impatience and by decreases in government spending, the growth rate

of population, the capital share, and the intertemporal elasticity of substitution.[14] A decrease in the depreciation rate has a theoretically ambiguous effect but, for reasonable parameter values, increases the consumption share of output. Similarly, an increase in the growth rate of productivity has an ambiguous effect but, for reasonable parameter values, increases the consumption share of output.

Of the large number of possible factors that can increase the consumption share according to equation (3.3), this paper focuses on several that are extant in the literature or suggested by recently observed changes in the economy. First, the share of output consumed by the government has declined over the last twenty years. A declining rate of government spending causes a consumption boom. Second, household wealth has increased despite low active saving, as documented in Figures 1 and 2. An increase in the capital stock causes a transitory consumption boom.

Finally, an increase in the discount rate of the representative agent increases the consumption share of output. While this cannot be observed directly, several extant theories imply an increase in the effective discount rate of the aggregate consumer.[15] First, the Social Security system is currently making large transfers from future generations to those alive today. Considering the representative agent derived from a life-cycle model, this increased intergenerational redistribution temporarily increases the effective discount rate of the representative agent.[16] Thus the Social Security system is considered as a potential explanation for the increase in the consumption share. Second, since households at different ages have different propensities to consume out of total resources, changes in the age distribution of the population change the effective discount rate of the representative agent. The aging of the baby boom generation and the increased life-span of the typical American have changed the demographic structure of the *U.S.* and may also have driven up the consumption share of output.

Third, in a model in which some households face large idiosyncratic risk or liquidity constraints, some saving is driven by precautionary or liquidity concerns. In the last twenty years, there have been an increase in the financial instruments employed by Americans and a significant increase in the ratio of debt to income. Thus I consider relaxed liquidity constraints as a possible explanation for the increase in the consumption to income ratio.[17] Finally, while not observed, there has been speculation that saving behavior differs by cohort. One version is that households who did not live through the Great Depression have a lower propensity to save than those who did. I examine whether there is evidence of an increase in

[14] For all of these effects, the change in consumption share at impact is the same as in the long run except that a decrease in the capital share can cause the consumption share to decline on impact.

[15] I do not consider one potential explanation advanced in Carroll and Weil (1994) and Paxson (1996). Habit formation tends to lead the growth rate of consumption to decline slowly following a slowdown in growth.

[16] That is, in a certainty model, intergenerational transfers to the present increase the propensity of the representative agent to consume out of current and expected output.

[17] See Caballero (1991), Ayagari (1993) and Carroll (1997).

the discount rate of the representative agent due to more patient older generations being replaced by more impatient younger ones.

4. A Quick Tour of Aggregate Evidence

In this section, I analyze which if any of the explanations just discussed are consistent with the observed changes in the aggregate economy. I focus on timing, on relative magnitudes, and on the composition of aggregate consumption. This first pass at the data is complemented by a thorough evaluation using household-level survey data later in the paper.

Before seeking to explain the increase in the consumption share, this section dismisses the possibility that consumption expenditures have increased while consumption has not. Suppose that there were a relative preference shift or price decline so that the representative household sought to increase the share of its consumption flows that are due to durable goods. Since the $NIPA$ measure expenditures rather than consumption, an increase in the share of consumption coming from the service flows from durable goods would generate a boom in consumption expenditures. In fact, however, the observed increase in consumption expenditures relative to income would not represent a decline in saving rates, but rather a shift of saving from capital to durable consumption goods.[18]

This supposition is easily rejected by an examination of household budget shares. Working with reference to GDP rather than total consumption, the ratio of expenditures on durable goods to GDP has remained steady since 1959, falling by a tenth of a percentage point from 1959 to 1979 and rising by a tenth of a percentage point since.

Turning now to the main explanations proposed in the previous section, there is little aggregate evidence that declines in government spending or appreciation of existing assets caused the increase in consumption to income. Since during the past twenty years the real interest rate was relatively high and the growth rate of consumption relatively low, the data suggest that the effective discount rate of the representative agent has increased.

4.1. Reductions in Government Purchases

Is the consumption boom driven by a decreasing share of output purchased by the government, due to the so-called "peace dividend" for example? In steady state, the canonical model of section 3 implies that the share of national output consumed by households and the government together is constant (equation (3.3)).

[18] Durable goods do not include housing. Housing services are counted as consumption, while housing stock depreciation and investment are counted as capital consumption allowance and investment. Changes in household wealth due to changing home ownership patterns are correctly reflected in the figures on saving. See Bureau of Economic Analysis (1987) and Bureau of Economic Analysis (1997).

A steady-state explanation can be quickly dismissed. Figure 3 shows that the share of output devoted to the purchases of both households and governments has risen over the past twenty years. The purchases of goods and services by governments have fallen by about 3 percentage points of GDP over the past ten years, but this decline is concentrated after most of the increase in the consumption share.[19] A second piece of evidence is provided by the real interest rate. Equation (3.2) shows that the real interest rate should be lowered by a decrease in the demand for output by the government. During the last twenty years, the real interest rate has been significantly higher than in the previous twenty.

There is however the possibility that a non-steady state explanation could work. That is, could the consumption boom be due to the expectation of both the currently observed decline in government spending and further declines in government spending in the future? This hypothesis is consistent with a high real interest rate and a high consumption share of output. If households expect lower government purchases in the future, consumption of the extra output available is smoothed by reducing investment and the capital stock in the present, thus increasing the real interest rate.

To evaluate this explanation, I ask what changes would have to be expected to rationalize the observed consumption boom. To keep things transparent, general, and easily reproducible, the present values are calculated holding the real interest rate constant. Such experiments provide a lower bound on the expected future declines in government spending.[20]

First, what is the expected steady-state share of government spending? The average ratio of government and consumption purchases to GDP from 1959 to 1979 is 84 percent. Given the current ratio of personal consumption expenditures to GDP of 68 percent, equation (3.3) implies that the expected steady-state ratio of government spending to output is 16 percent.

Second, what accumulated value from the consumption boom must be recovered from lower government spending? Consider first the counterfactual that the consumption ratio remained at its $1959-1979$ value over the $1979-1998$ period. The present value of the excess of the observed consumption series against this alternative stands at $5,782$ billion 1992 dollars, or three-quarters of a year of GDP, when accumulated at a three percent real interest rate. The decline in government

[19] Also, the constant consumption share and the declining share of government purchases over the past five years suggests no "crowding in" of consumption in response to the reduction in the share of government expenditures.

[20] The fact that the partial equilibrium experiment provides a lower bound can most easily be seen in two steps. First, consider the household budget constraint. Because the capital stock declines as consumption rises and then rises as government spending further declines, the real interest rate is high when the household is borrowing from the future (reducing capital below the steady-state level). Thus, to "pay off" the early consumption boom requires greater saving (a greater decline in government spending in the future) than if the interest rate had been constant. Second, since we see the current ratio of consumption to income and the real interest rate should decline as we get to steady state, the steady-state consumption to income ratio is actually higher than the observed one, thus requiring a still-lower steady-state level of government spending.

spending as a share of GDP since its local peak of 21% in 1987 cumulates to only $1,400$ billion 1992 dollars to date.

One path of government purchases that can rationalize the consumption boom, is if the ratio of government purchases to GDP declines a half a percentage point a year to 13 percent, stays there for 15 years, and rises again at a half a percentage point per year to 16 percent.

Thus, to rationalize the consumption boom from this source requires extreme expectations of declines in government purchases. Additionally, we have moved from the observed to the unobserved. To date, no government spending movements have occurred that can rationalize more than a small fraction of the consumption boom.

4.2. Appreciation of Existing Assets

As shown in Figure 2*b*, the value of assets owned by the representative household have been increasing relative to its income. Can this explain the increase in the consumption to income ratio?

First, what might generate the large increases in the ratio of net worth to income while investment share is low and the real interest rate is high? If households realized that the capital stock were higher than they had thought, then the consumption share would increase, but, counterfactually, the real interest rate would be low. Instead,suppose that households expect a big increase in output in the future. If so, then households deccumulate capital, the real interest rate rises and consumption rises as a share of output. These real interest rate, consumption and output movements are as observed in the data. If in addition, firms must invest now, for example in information technologies, in order to reap these future productivity gains, then it is also possible for the theory to predict an increase in net worth to income, as in Greenwood and Yorukoglu (1996). If this investment is not measured as output or investment, then consumption rises as a share of output.[21]

An alternative theory is simply that asset prices follow fads or bubbles. In either case, two problems are encountered in trying to explain the consumption boom with the increase in wealth.

First, the timing is wrong. The increase in the wealth to income ratio is mainly due to the increases in financial assets — largely the stock market — as shown in the lowest line in Figure 2*b*. This increase occurs primarily in the last five years, a time when personal saving is declining but the consumption to output ratio is constant.

Second, focussing on the years over which the consumption share of income increased, the total increase in the ratio of net worth to income from the late 1970's to the mid-1990's is about one-third of a year of GDP. The marginal propensity to consume out of wealth must be one-sixth to rationalize the consumption boom.

[21]See Greenwood and Yorukoglu (1996) and Greenwood, Hercowitz and Krusell (1997) or assume that the future increase in productivity is associated with certain existing pieces of capital. Another possible shift in technology is a decrease in the capital share of the economy.

If one chooses such a high marginal propensity to consume, however, one is then puzzled as to the lack of consumption response to recent increases in wealth.

The role of wealth accumulation becomes more plausible if one ignores timing and argues instead that through budget constraints consumption must be related to wealth. The increase in net worth to income from the late 1970's to 1997 amounts to two-thirds of a year of GDP. The marginal propensity to consume out of wealth need now only be 9 percent to rationalize the consumption boom.[22]

Thus, while the aggregate data cast some doubt on the role of wealth, this explanation for the consumption boom is a main focus of the subsequent analysis of household data.

4.3. Increases in Impatience or the Propensity to Consume

As noted at the end of Section 3, several extant explanations argue that the effective discount rate or the propensity to consume of the representative agent has increased. Such an increase is consistent with two main coincident facts. First, as already mentioned, the real interest rate is high during the consumption boom relative to the previous two decades. This suggests that the demand for output is relatively high. Second, as documented in the first two rows of Table 1, the growth rate of real consumption per capita actually has slowed. Within the context of a Ramsey economy, the Euler equation governs consumption growth. Without a change in the effective discount rate, a higher real interest rate should be associated with a faster average growth rate of consumption, not a slower one.

I now turn to two of the explanations discussed at the end of section 3: increases in government transfers from future to present generations and financial innovation and increases in debt. These explanation are also evaluated in sections 6, 7 and 8, using cross-sectional implications of these theories and household data.

4.3.1. Increasing Government Transfers to Older Generations

During the period of the increasing consumption share of output, the U.S. government has increased its reallocation of wealth from future to current generations. In a pure life-cycle model, the beneficiaries of these transfers consume more than their pre-transfer wealth, while other generations consume less. In the U.S., Social Security and Medicare are the largest of these programs, and the payments to the elderly have been consistently rising as has the composition of medical care in total consumption. Gokhale, Kotlikoff and Sabelhaus (1996) argue that this redistribution can explain the consumption boom.

To provide a first evaluation of this explanation, the second panel of Table 1 presents the budget shares of different categories of consumption, including medical care. The boom in consumption is more than entirely due to increased consumption of services, of which medical care is a major component. The output share

[22] Poterba and Samwick (1995) and Ludvigson and Steindel (1999) also demonstrate that the high-frequency relationship between stock market value and consumption is weak.

of purchases of goods— nondurable and durable — has declined by nearly 10 percentage points since 1959. Two thirds of this decline is a steady decrease in the share of consumption that is food. Within services, the largest increases in consumption are due to spending on medical care and other services.[23] From 1979 to 1998, the growth in the share of medical care is 4.3 percent and the growth in the share of other services is 4.2 percent, both large when compared to the 5.9 percent increase in the total consumption to GDP ratio. This seems to suggest that the consumption boom can largely be explained by government provision of medical care free-of-charge to the elderly.

However, a slightly different picture emerges if one compares these recent changes, that occur contemporaneously with the consumption boom, with the changes that have occurred over the previous twenty years, when the consumption to GDP ratio was constant. Table 1 again reveals that services growth is, at least in an accounting sense, the cause of the recent consumption boom. But, relative to growth over the previous twenty years, the increased consumption of services is more evenly distributed among non-medical non-housing services, medical services and housing. The change in medical care services as a share of output from 1979 to 1998 exceeds the change over the previous twenty years by 1.4 percentage points. This is consistent with the fact that the transition to Medicare is largely completed prior to the consumption boom.[24]

In sum, there is evidence that the consumption boom is concentrated in spending on services, but not that this increased spending on services is disproportionately concentrated on medical care. Since this evidence is far from conclusive, I later evaluate the role of intergenerational transfers, including Medicaid, by studying which households are "overconsuming" relative to their ages, wealth, and incomes, and asking whether these households are in cohorts that are receiving large intergenerational transfers. In good macroeconomic tradition, the remainder of the paper will focus on output as one good.

4.3.2. Financial Innovations and Increases in Debt

During the past twenty years, gross debt has risen as a share of disposable income, increasing from 0.7 to nearly 1, shown in Figure 2*b* as the difference between total assets and net worth. If this increase represents relaxed liquidity constraints or financial innovation that allows households that might previously have been constrained to borrow to support consumption, then this innovation would lead to a transitory consumption boom. During the past twenty years, credit cards have become more widely available and an increasing amount of debt is held on these

[23] Other services include transportation services and household operations (which are usually their own categories) and miscellaneous services related to: clothing, accessories and jewelry such as cleaning, repair, and storage; personal business such as banking, legal, and funeral services; recreational services such as cable TV, club memberships, theatre tickets, and pet-related costs; religious activities; foreign travel; and finally education and other day-care costs. See Bureau*d*of Economic Analysis (1990).

[24] See Bosworth (1996).

cards. Also, the minimum down payment required to purchase a house has declined and the number and prevalence of financial instruments available to reduce home equity have increased.

Financial innovations are not able to account for a large increase in consumption. As noted in the previous subsection, the share of expenditures on housing services rises during the consumption boom. However, the increase in the ratio of debt to output is just over twenty percent. As calculated in section 4.1, the present value of the consumption boom is three-quarters of one year of GDP. If the increase in the debt ratio were entirely caused by an exogenous increase in households' ability to borrow, then financial innovation could explain at most thirty percent of the increase in consumption to income to date.

4.4. The Limits of Aggregate Evidence

Using only aggregate data, a significant difficulty in understanding the decline in the saving rate is lack of exogeneity. Thus the paper now turns to household-level data. This has three advantages. First, the composition of households has changed significantly over the past twenty years. There are more retirees, more single parent families, and greater dispersion in household income. The paper uses household-level information to evaluate whether such changes have caused the decline in the saving rate. Second, several possible causes of the consumption boom give strong predictions about the cross-sectional distribution of consumption ratios. For example, intergenerational transfers are expected to be positive for the currently elderly and negative for the currently young. Finally, absent full consumption insurance, household propensities to consume out of idiosyncratic asset values and income levels can be used to estimate the response of the aggregate economy to these variables.

The next section describes the construction of a novel dataset that combines information from two household-level survey datasets and $NIPA$ data to generate a panel dataset with information on consumption, income and wealth at the household level. The remainder of the paper uses this dataset to evaluate theories of the increase in aggregate consumption relative to income.

5. Constructing a Household-Level Dataset

In order to study the consumption behavior of households, I employ the Consumer Expenditure Survey (CEX) to impute the consumption of services and nondurable goods to each household in the Panel Study of Income Dynamics ($PSID$) from 1979 to 1994. This yields a panel dataset on consumption of households that includes a large set of demographic and income information as well as three years of detailed wealth information. Additional details are provided in the appendix.

5.1. The PSID

The $PSID$ has been used extensively to study year-to-year fluctuations in consumption and the main characteristics of the dataset are reasonably well known. For the present analysis, using $PSID$ households as the unit of analysis has three main advantages. First, the survey provides panel data over much of the time period of interest on over $5,000$ households per year. Data from 1979 to 1994 are used to match the timing of the consumption boom.[25] The $PSID$ provides weights so that the means in any year or category of household can be aggregated to produce a nationally representative sample.

Second, the Survey has repeated measures of food consumption and excellent information on household income. The main measure of consumption is usual weekly food consumption, and this information has been gathered in every year of interest except 1988 and 1989. Food consumption is measured with error, and this has hampered studies working with Euler equations and relating annual consumption changes to observable variables. In much of this study, the focus of interest will be long term movements or movements across groups of people so that this mismeasurement creates fewer difficulties. The fact that food consumption is not typical of all consumption expenditures is more of a concern, and this concern leads to the joint use of the CEX, as subsequently described.

Income in the $PSID$ is total post-transfer, pre-tax income, so that it is not completely comparable to National Income in the $NIPA$. Nonetheless, as demonstrated in the appendix, the ratio of food consumption to income constructed from aggregating the $PSID$ data has the same temporal pattern as that of the $NIPA$. The correlation between the $PSID$ series and the $NIPA$ series is 0.93.

Third, and most importantly, the survey contains accurate information on wealth holding of households in 1984, 1989, and 1994, a time period covering the heart of the consumption boom. Such information is not available in the CEX alone. The $PSID$ data on wealth include wealth held in saving and checking accounts, money market accounts, certificates of deposit, bonds, stocks, mutual funds, IRA's, cash value of life insurance, trusts and estates, main home, second homes, investment real estate, cars, trucks, boats, motor homes, farm and business wealth, and collections of things for investment purposes (e.g. baseball cards), all less credit card, mortgage, and "other" debts. The wealth data are comprehensive and do an excellent job of reproducing the wealth of the bottom 99% of the wealth distribution in these categories of wealth (Juster, Smith and Stafford (forthcoming)). The only real shortcoming of the $PSID$ wealth data is that pension wealth is unavailable. The $PSID$ does report whether the household has a pension and this information is used here.

In order for a household to be included in the analysis, it must have all the necessary information for the year in question. Further the observation is dropped if any of the necessary information is a major assignment made by the $PSID$ staff.

[25] These include income information from the Survey year 1995. The 1994 and 1995 data are in early release form and thus the relevant variables must be constructed from raw data.

5.2. The CEX

In order to use the $PSID$ to analyze the increase in the consumption share of output, this work imputes the consumption of nondurable goods and services for each household. I first estimate the relationship between this larger measure of consumption and a household's level of food consumption and demographic characteristics using data from the Consumer Expenditure Survey (CEX). The consumption of nondurable goods and services of households in the $PSID$ is then predicted using this estimated relationship.

The CEX is conducted by the Bureau of Labor Statistics in order to construct baskets of goods for use in the bases for the Consumer Price Index, and has been run continuously since 1980. The survey has excellent coverage of consumption expenditures, reasonable data on liquid assets, and income information of moderate quality.[26] The survey interviews about $5,500$ households each quarter and has households keep records of consumption expenditures that are then collected by the survey at the end of four three-month interview periods. About half of all households make it through all the interviews and sample weights are given so that a representative sample of non-rural households can be recovered. The CEX represents the best source of information on household consumption across a large set of categories.

The data used here come from the Family files of the CEX from 1980 to 1993 and from extracts made publicly available by the Congressional Budget Office and John Sabelhaus through the $NBER$.[27] Each household contributes one data point to the employed sample. I drop any household that is classified as an incomplete income reporter, that has any of the crucial variables missing, or that does not report an income measure contemporaneous with the consumption data.[28] I construct variables measuring food consumption and consumption on all $NIPA$ categories of nondurable goods and services consumption. Income is pre-tax total family income to match the concept in the $PSID$.

Finally, as for the $PSID$, I construct the ratio of food consumption to income from the CEX and compare this to the $NIPA$ series. The correlation between the CEX and $NIPA$ series is 0.78, which is not as high as that from the $PSID$, but, as discussed in the Appendix, an acceptable level for the present purposes.

5.3. Imputing Expenditures on Nondurable Goods and Services

Turning to the imputation of consumption for households in the $PSID$, two important factors drive the specification of the imputation. First, what are the correct theoretical concepts that shift the relative utility of consumption of food and nonfood items? Given that food has declined significantly as a share of consumption over the period of interest, to impute nondurable and services consumption to

[26] See Lusardi (1996), Attanasio (1994), and Branch (1994).

[27] See Bureau of Labor Statistics (1980-1993) and http://www.nber.org/ces_cbo.html.

[28] This procedure cuts nearly all households that are listed in the CBO/Sabelhaus/NBER data as not completing all the interviews. The weights adjusted by the CBO for attrition are employed.

households it will be necessary both to recognize that the relative price of food changes through time and that food is a necessity, so that its budget share declines with wealth. Further, household characteristics such as family size, the number of earners, and retirement status may shift the relative utility of food consumption versus consumption of other goods.

Second, what variables are measured in similar ways in both surveys? The imputation is only valid if the regressors used in the estimating equation are the same variables as those in the predicting equation. As discussed in the appendix, there is some variation in the relative levels of the consumption and income series, but the factors of interest are the time trends. For all the regressors, the survey questions, the levels, and the time trends are compared between the surveys and they match reasonably well.

The imputation proceeds in four steps. First, using the CEX data, the log of expenditures on nondurable goods and services is regressed on a cubic polynomial in the log of food consumption and a set of regressors designed to allow preferences for relative consumption to vary by family size, age, education level, labor force status, and retirement status. To capture differences in relative prices of goods over time, the mean is allowed to vary by year. The regression employs $37,730$ households and explains 80 percent of the variation in household consumption.

Second, the estimated parameters are used to predict consumption of nondurable goods and services for each household in the $PSID$. Third, the imputed consumption for each household is treated as a relative consumption level and the total consumption across households is scaled up to include medical purchases by the government. This step is similar to that of Gokhale et al. (1996) who assign medical consumption across ages. Their medical care adjustments employ more detailed age-specific adjustments but do not assign these expenditures in relation to individual consumption. Finally, the consumption of nondurable goods and services in the $NIPA$ in each year is allocated across households in proportion to each household's consumption from the third step.

After this imputation, I have a true panel dataset that covers 16 years from 1979 to 1994 and contains measures of income, nondurable and service consumption (for all years except 1988 and 1989), and wealth in 1984, 1989 and 1994. I turn now to describing the evolution of consumption ratios across broad groups of the population. All nominal data are made real using a price index constructed by dividing nominal consumption of services and nondurable goods by the same real quantity, where nondurable and services consumption is made real using the $NIPA$ chained price indexes. Data for the second quarter of the year of interest are used.

From here on the term "consumption" is used interchangeably with the more cumbersome term "consumption of nondurable goods and services."

6. Growth and Demographic Structure: Age Cohort and Time Effects in Consumption

The US has experienced a large increase in the share of the population that is over 65 years of age and a bulge in the population distribution associated with the aging of the baby boom generation.[29] If households of different ages have different propensities to consume out of lifetime income, then there is variation in the representative agent's discount rate. For example, middle-aged households wish to consume at a greater rate than young or old household since they tend to have more members. Thus an economy in which a population bulge is entering middle age looks like a canonical Ramsey economy with a temporarily higher effective discount rate. Are the observed fluctuations in the U.S. age distribution leading to fluctuations in the discount rate of the representative agent that are in turn pushing up the consumption to income ratio?

6.1. Who Are Consuming More of Their Income?

Over the period in question, the elderly as a group have increased their share of annual consumption. This fact suggests an important role of decreasing lifetime wealth of the young and/or increasing transfers to the elderly. However, this trend significantly predates the current data and the consumption boom.[30] Following in the footsteps of previous studies using micro data, the analysis of the household data begins by describing the evolution of consumption and consumption ratios across different age groups and time periods.[31] Next, this section uses a simple life-cycle framework to identify the role of demographics in the consumption boom.

The analysis first groups the data into birth cohorts and age groups. Table 2 shows the cells and the cell sizes chosen for the analysis. Ages are grouped into 13 five-year cells and the cohorts are also split into 12 cells.[32] The number of households in each cohort cell and age group varies over time. There are 14 years of data, spread over the 16 years 1979 to 1994. In general, each cohort group and age group will be denoted by the middle age or year in its range. The amount of data afforded by the $PSID$ is a significant advantage: there are over $80,000$ observations on household consumption and income. This is an unweighted look at the data. Sample weights imply quite a different age and cohort distribution of the data, one that is representative of the U.S. noninstitutional population.

To begin, I use the sample weights and data in each cohort group to construct

[29] In addition, there is a long-term trend towards slower population growth in the US.

[30] See Gokhale et al. (1996) Figure 1.

[31] This approach is employed in the study of consumption and saving by Deaton and Paxson (1994), Attanasio (1997a), Deaton and Paxson (1997), and Alessie, Kapteyn and Lusardi (1998).

[32] Two "partial" cohorts are in the sample for too little time to properly identify their actual cohort effect. Of these partial cohorts, the youngest cohort is only observed in the relevant age range for about half the sample. The oldest has some members in the sample in every year, but fewer than 50 in each year of the 1990's. These partial cohorts are used only in a subset of the analysis, and when this is done it is noted.

a measure of the average log consumption of each cohort at each age. Figure 4*a* displays the consumption of each cohort at different ages. The life-cycle pattern of hump-shaped consumption is clearly visible.[33] Also noticeable is the artificially sudden rise in consumption that occurs at age 65 due to the allocation of medical expenditures by the imputation procedure.[34] The figure shows, for any cohort, the combination of both age and time effects at work. None of the effects are separately identified. It could be that all households have the same lifetime wealth and that the "endpoints" of each segment do not join due to time effects that raise the endpoint of each cohort's age-series of consumption. However, productivity growth implies that younger cohorts are richer and so consume more than their elders did at the same age. If there were no time effects, then consumption profiles of younger cohorts would lie above those of their elders solely due to cohort effects that would be due in turn to productivity growth. To identify the separate impacts of age, time and cohort requires identifying assumptions, which are provided shortly.

Figure 4*b* displays the same set of information as Figure 4*a*, but by year and for only four cohorts. This figure shows that over the period of the consumption boom, the cohort whose consumption has risen the most is that of households born between 1955 and 1959, the youngest cohort. While this would seem like evidence that this younger group is, in an accounting sense, the cause of the consumption boom, in fact, the age-profile of consumption for this cohort should be increasing.

Figures 5*a* and 5*b* show the total consumption of each cohort divided by the total income received by that cohort by age and time respectively. Figure 5*a* emphasizes the clear life-cycle pattern of consumption ratios, in which the young save and the elderly dissave. Again these patterns are confounded by the inability to see people of different cohorts at the same age and in the same year. Looking at the general shape of the profile, one sees a mixture of effects at work. That is, since the profiles for different cohorts nearly join neatly or overlap when observed at the same ages, it may look as if the effect of cohort on saving behavior is small. In fact, however, these profiles may not join or overlap if time effects are removed. Figure 5*b* displays the combination of the effects of age and time on each cohort of households. The cohort born between 1925 and 1929 clearly has the sharpest rise in consumption ratio over the period; however, the same caveat that applies to the increasing consumption of the young applies here. During the 16 year period examined, the youngest households in this cohort age from 50 to 65 and the oldest households in this cohort move from 54 to 69. Thus, life cycle considerations suggest that this group should move from saving to dissaving.

Finally, Figures 6*a* and 6*b* display the profiles of the average of household-level consumption growth. As is typical of household data, the growth rates of consump-

[33] This pattern has many interpretations and has been the subject of much debate, see for example Carroll and Summers (1991), Attanasio and Browning (1995), and Gourinchas and Parker (1997).

[34] In the analysis of consumption levels, this feature of the imputation only biases the estimated age effects. In the growth-rate regressions, the artificial consumption growth over these years is removed by a dummy variable.

tion display a fair amount of variation, but the life-cycle figure still captures a broad age-pattern in the same way as figures $4a$ and $5a$.[35] Household-level variation is potentially useful for identification of the underlying causes of the consumption boom. The profiles by time seem to have more measurement error, although the data do pick up the aggregate growth following the 1982 recession and the decline in consumption growth in the 1991 recession.

6.2. Identifying the Impact of Demographics

In this section, each household's consumption and income is decomposed into a portion specific to the time period, a portion specific to its birth cohort, a portion specific to its age, and a final portion specific to the individual household. By making the household-specific portion by definition mean zero for each age, cohort, and time grouping, the aggregate consumption ratio can be reconstructed from a weighted combination of age, time, and cohort components for each time period. Separately identifying age, cohort and time effects requires an identifying assumption.[36]

The canonical methodology for separately identifying the effects of age, cohort, and time in saving rate data is to assume either that time effects are unimportant or that they are mean zero and orthogonal to a time trend (Attanasio (1997a) and Deaton and Paxson (1994)). Income and consumption are composed of four additive effects: a time effect specific to the year the household is observed; a cohort effect that captures permanent differences in wealth and situation; an age effect, that captures the typical household's saving profile over their life; and finally a household-specific component, uncorrelated with the first three. In the absence of fluctuations, the stripped down life-cycle model of Modigliani and Brumberg (1956) predicts identical age profiles for each generation and cohort effects that depend on lifetime resources. Attanasio (1997a) and Paxson (1996) provide evidence that age-profiles over long time horizons conform reasonably well to this model.

I assume that the time effects are mean zero and orthogonal to a linear time trend. The consumption increase can then only be traced to differential saving behavior of different generations or to different shares of the population at different ages. While this decomposition is informative without a direct structural interpretation, a simple life-cycle model predicts these effects. In the basic life-cycle model, the household consumption ratio, $\frac{C}{Y}$, can be written as the marginal propensity to consume at that age times the household's wealth:

$$\frac{C_h}{Y_h} = MPC_a \left(\frac{NPVY_h + W_h + NT_h}{Y_h} \right) \tag{6.1}$$

[35] An alternative approach would be to average consumption by year and group first, and then first difference. But the amount of noise in household-level consumption growth does not seem to be sufficient to require that one look only at consumption growth by group.

[36] Smoothing the data using age and cohort groups can provide an artificial identification. To avoid this, all members of a cohort are assigned to the same age so that age=year-cohort and the identification of the linear relationship among the effects requires an identifying assumption. The results, once identification is imposed, are substantively unchanged by this modification.

where $NPVY$ denotes the net present value of human wealth, MPC denotes the propensity to consume out of total resources, h denotes the household in question, a denotes age, W denotes wealth, NT denotes the present value of net transfers. All wealth measures are as of the start of life. The marginal propensity to consume out of wealth is allowed to vary by age, presumably due to changing family size, time until death, and possibly due to changing preferences (or even unmodelled precautionary saving).

Taking logs yields

$$\ln\left(\frac{C_h}{Y_h}\right) = \ln MPC_a + \ln\left(NPVY_h + W_h + NT_h\right) - \ln(Y_{h,t})$$

and so that aggregate consumption to income ratio can be exactly reconstructed after the decomposition, the approximation $\ln\left(\frac{C_h}{Y_h}\right) \approx -\frac{S_h}{Y_h} = \frac{C_h}{Y_h} - 1$ leading to

$$\frac{C_h}{Y_h} = A_h + B_h + T_h + \varepsilon_h \tag{6.2}$$

where $A_h \equiv 1 + \ln MPC_a$ plus the sample average of $\frac{C}{Y}$, B_h is the average of $\ln\left(NPVY + W + NT\right)$ across households in the same cohort as h less the sample average, T_h is the average of $-\ln(Y_{h,t})$ across households in the same year as h less the sample average, and finally ε_h is that share of the consumption ratio not explained by the three effects. Under certainty, the cohort effect depends only on lifetime resources. Fluctuations in income deliver time effects. Note that in estimation, sampling error falls naturally into a time effect.[37]

Before decomposing the ratio of consumption to income as shown in equation (6.2), I decompose household consumption into age, cohort and time effects. Household consumption is regressed on a complete set of age dummies, a set of time dummies less two, and a complete set of cohort dummies less one. Cell weights are used in the regressions so that the relative importance of a given cell in generating the aggregate is accounted for. The regression constrains the coefficients on the time dummies to sum to zero and to be orthogonal to a time trend. The coefficients on the cohort dummies are constrained to be mean zero.

Figure 7*a* shows the decomposition of household-level consumption. The age profile of consumption rises with age and declines less than the rough profile of figure 4. This difference is due to the cohort effects that steadily increase over the century. Each successive cohort consumes more, presumably because its lifetime resources are greater.[38]

Figure 7*b* shows the same decomposition applied to the consumption ratio, as in equation (6.2). Consumption and income are separately constructed for each cell of cohort, age, and year and the consumption rate is constructed for each

[37] The extant set of models that yields time, age, cohort decompositions maintain the dual assumptions of certainty and a constant real interest rate.

[38] The rate of increase of the cohort effects clearly slows over time, consistent with the slowing of productivity growth.

cell by dividing total consumption by income.[39] The age effects in consumption ratios show a typical profile of nondurable and service consumption rates for any generation. Households during their working lives consume less than their incomes, and a roughly constant fraction as they age. As income declines at age 60 and during retirement, households consume significantly more than their incomes. The implied saving profile looks quite similar to the predictions of the textbook life-cycle model.

Turning next to the cohort effects, there is clear evidence that the younger cohorts are bigger spenders than the older cohorts, relative to their incomes. The effect is large, with the cohorts born most recently on average consuming over fifteen percent more of their income than the oldest households. What causes such large differences? Within the framework of the simple life-cycle model above, this higher level of consumption comes from younger cohorts having higher wealth relative to income, such as from net government transfers or bequests.

The role of increases in wealth will be evaluated shortly. The role of intergenerational transfers is studied closely by Gokhale et al. (1996). Gokhale et al. (1996) construct certainty-equivalent wealth levels in a life-cycle model and examine saving rates from 1963 to 1989. Their decomposition blames the declining national saving on government transfers to households that are elderly by 1989. If the pattern observed in Figure 7*b* were due only to intergenerational transfers, the net transfers to the youngest cohorts would have to be larger than those to the older cohorts. This is somewhat implausible and inconsistent with the intergenerational transfer distributions constructed by Gokhale et al. (1996).

In sum, within the context of a basic life-cycle model, fiscal transfers across generations alone cannot explain the consumption boom. Transfers may be leading today's elderly to consume a larger share of their incomes than the elderly of two decades ago. But Social Security cannot explain the propensity of cohorts born more recently to consume a higher fraction of their incomes than the current elderly.

6.3. Can Changing Demographics Explain the Consumption Boom?

This subsection demonstrates that the changing distribution across age groups in the U.S. does not explain the increase in the ratio of consumption to income. According to the decomposition of section 6.2, there are two possible explanations of the decline in saving. First, the weight given to different age effects changes as the shares of different age groups in the population change. For example, as the elderly have become an increasing share of the population, they may have pushed the aggregate saving rate down because the elderly consume a larger fraction of their incomes than other age groups. Second, the cohorts that are higher consumers may move to the ages at which their consumption and incomes are higher and so push up the aggregate consumption rate. Lower consumption cohorts may also die and be replaced by higher consumption cohorts.

[39] Similar conclusions are reached employing separate identification of effects in consumption and income at the household level.

This subsection uses the estimated effects to consider partial-equilibrium alternative scenarios in which different weights are given to different effects in generating the aggregate consumption ratio. The aggregate consumption to income ratio for each year, denoted $\widehat{\left(\frac{C}{Y}\right)}_t$, can be reconstructed as

$$\widehat{\left(\frac{C}{Y}\right)}_t \equiv \frac{\sum_{i \in I_t} w_i \hat{B}_i}{\sum_{i \in I_t} w_i} + \hat{T}_t + \frac{\sum_{i \in I_t} w_i \hat{A}_i}{\sum_{i \in I_t} w_i}$$

where i indexes age-cohort-year cells, I_t is the set of cells for which the year is equal to t, w_i is the population weight associated with that cell, $\hat{B}_i$ is the estimated birth-year or cohort effect, $\hat{T}_i$ is the estimated year effect, and $\hat{A}_i$ is the estimated age effect.

Figure 8*a* displays the reconstructed consumption to income ratio without cohort effects.[40] Figure 8*b* shows the consumption ratio with age and time removed, leaving only the effect of cohorts aging. These figures show that the consumption boom is not due to the changing age distribution. Instead, the decline in saving occurred because each successive generation consumed more of its income than the previous generation at that age.[41]

This conclusion matches the general consensus of research in this area that the age distribution of the population has little effect on national saving (Bosworth, Burtless and Sabelhaus (1991), Paxson (1996), Attanasio (1997a) and Deaton (1997)). All of these papers employs slightly different methodologies and data, and all blame cohort rather than age effects for saving rates. Attanasio (1997a) finds that those born between 1925 and 1939 account for an unusually high share of national consumption. Gokhale et al. (1996) attribute the decline in saving between 1960 and 1990 to the large share of resources flowing from future generations to the generation that is currently elderly.[42] The findings of the remaining sections concur that age dynamics have little to no impact on the consumption ratio.

The balance of this paper is devoted to a fuller investigation of the structural interpretation of these all-important cohort effects. In this section, the cohort effects represent differences in lifetime resources because the environment is assumed so simple that no other explanations are present to compete. There are two reasons to be sceptical of such a simple interpretation. First, the observed pattern of

[40] This analysis is conducted including the "partial" cohorts so as to replicate the aggregate time series. This reconstructed consumption ratio has a slightly lower increase over the period than the raw data, which implies that the true cohort effects for the extremely old and young are larger in absolute value than the endpoints that are used for them. Also, the changing numbers of these households over time induce some year-to-year fluctuations in the reconstructed ratio that are not due to time effects.

[41] The same conclusion and similar pictures are obtained if I instead separately remove cohort effects from consumption and income at the household-level and reconstruct time series without cohort effects in either series.

[42] About half of the increase in consumption they attribute to an increasing propensity of the elderly to consume, a propensity that is not identified as due to age, cohort, or time.

fiscal transfers is not consistent with the estimated pattern of the cohort effects. Second, there are important observed changes in the U.S. economy that call into question the simple identification scheme of this section. Differences in real interest rates, shocks to wealth, and different rates of time preference across generations all invalidate the identification assumptions employed here by altering the age profile of consumption across households.

To address these shortcomings, the next section augments the simple life-cycle decomposition. I allow for uncertainty and model the cohort effect as due both to the permanent component of income and to wealth holdings. Estimating, a linear approximation to the household consumption policy function, I again find that the appreciation of assets alone cannot explain the consumption boom.

7. The Role of Wealth

This section considers a realistic but simple model of household behavior and estimates an approximate consumption policy function for each household.[43] The procedure of this section does not assume that time effects are mean zero or that the agent's environment is certain. The consumption boom is sourced to the changing age distribution, time effects, and the changing distributions of wealth and the permanent component of income.

7.1. An Organizing Model

Each household in the economy chooses consumption to maximize expected lifetime utility

$$Max \quad E_s \left[\sum_{t=s}^{T} \beta^{t-s} \nu_a u(F_t, C_t) + \beta^{T+1-s} V^D_{T+1}(F_{T+1}, X_{T+1}) \right]$$

where E_s is the expectations operator conditional on all information available at time s; β is the discount factor; ν shifts utility as households age, F is a family-size adjustment that normalizes consumption to per-capita terms, X_t is household cash-on-hand; and $V^D(\cdot)$ captures the possible value of cash on hand remaining at death. Household choices are constrained by an intertemporal budget constraint that represents the evolution of liquid assets or cash on hand, X_t, and a liquidity constraint that they must maintain positive net wealth

$$\begin{aligned} X_{t+1} &= \tilde{R}_{t+1}(X_t - C_t) + (1-\tau)Y_{t+1} \\ X_t &\geq C_t \end{aligned}$$

where $\tilde{R}_{t+1}$ is the gross after-tax rate of return on the household's optimal portfolio, and Y_t is disposable non-asset, pre-tax income.

[43] Recent work that estimates consumption functions includes Carroll (1994) and Parker (1998).

The household bases its consumption upon its current state and its expectations about the future. That is, household consumption is described by an optimal policy function of the payoff-relevant state variables. In order to choose its current consumption level, the household needs to know its current and expected future resources, its family size, the time-horizon over which it is alive, and the possible investments and rates of return available to it. In order to forecast future income, I assume that the household only requires the permanent component of its income, P_h, the aggregate state, A_t, and its age.[44] I assume that the household requires only knowledge of the aggregate state to forecast optimally future rates of return.

Under these assumptions, the consumption function for household h can be written solely as a function of family size, wealth, income, age, the permanent component of income, and the aggregate state

$$C_h = F(\digamma_h, X_h, age_h, P_h, A_t). \tag{7.1}$$

Since different cohorts may still have different preferences for consumption above and beyond their state variables, and since there may be a role for different intergenerational transfers by cohort, the exclusion of birth year from the consumption function is tested.

7.2. Estimation Strategy

A log-linear approximation to the policy function is estimated of the form:

$$\ln C_h = g(\digamma_h) + h(X_h) + f(age_h) + \beta \ln P_h + T_t + \varepsilon_h \tag{7.2}$$

where the residual represents measurement error in the level of consumption, and T_t is a year effect that captures the aggregate state, that is, changing expectations about the future. This equation is estimated on the $PSID$ data in 1984, 1989, and 1994, the years in which, as previously discussed, the $PSID$ has an accurate reporting of household wealth. The data are constructed from the $PSID$ data already employed, with the addition of these three years' wealth supplements and the following two constructions.

First, I construct a measure of consumption in 1989, a year in which the $PSID$ does not report food consumption. Consumption from 1990 is used instead and deflated for each household by the aggregate growth in consumption between 1989 and 1990. Since any innovation to marginal utility between 1989 and 1990 should not be predictable by anything known in 1989, such as what is on the right hand side of the 1989 regression, this substitution should not adversely affect the results. Second, I construct the permanent component of income as the forecast of the log of current income from two lags of the log of family income, education and age group dummy variables. This forecast is done separately for retired and non-retired households.

[44] The permanent component will be defined shortly. I will also consider a case in which current income is necessary for predicting future income.

The function $g(F_h)$ consists of the size of the family and the number of children in the family. A set of dummy variables representing the five-year age groupings capture the age effects on consumption, $f(age_h)$. Finally, wealth is included in the regressions as the log of wealth if it is positive, a dummy for wealth being zero or negative, and a dummy for whether the household has pension.

How does this model differ from the age-time-cohort decomposition of the previous section? The key differences are two. First, the model includes directly both wealth and the permanent component of income in place of the cohort effect of the previous section. The behavior of consumption can then be traced to both this observable version of the cohort effect and to time effects and omitted elements of the cohort effect. Second, the time effects are not constrained to be orthogonal to a linear trend. Thus they can explain trend movements in consumption that are not explained by increases in wealth, the changing age distribution, and so forth.

Equation 7.2 is estimated on the entire sample of weighted data with imputed real nondurable and services consumption as the dependent variable. The time effects capture expectations, real interest rates, and all aggregate conditions. The only source of variation is cross-sectional. The goal of the exercise is to see whether the behavioral relationships estimated from household data can explain the consumption boom when time series variation is substituted for cross-sectional variation.

7.3. Behavioral Evidence on the Consumption Boom

Table 3 displays the results of estimation of four different specifications and the implied increases in the ratio of consumption to income due only to changes in the distribution of wealth to income over the period.[45]

The marginal propensity to consume out of wealth is estimated to be around four percent. As noted in section 3, ignoring timing, a marginal propensity to consume out of wealth of 9 percent can rationalize the entire twenty-year consumption boom. Over the 10 year period being studied here, however, wealth increased in relation to income only over the first five years; during the second five years the distribution of wealth spread out, so that the number of low wealth households increased despite no significant change in the mean wealth to income ratio.

The estimated relationship between consumption and wealth is not linear, in that the cluster of low-wealth households have more consumption than would be implied by the relationship between wealth and consumption for higher wealth households. The $PSID$ does not measure pension wealth, but the presence of a pension increases consumption by between two and a half and five percent.

When interpreting the income variables— the current income and the permanent-component or expected income— one must keep in mind that the time-effects remove mean long-run correlations. That is, if the model were identified from the time dimension, then rising incomes and consumption together with the budget

[45] See Hurst et al. (1998) for a detailed description and analysis of the distribution of wealth in the $PSID$. See also Sabelhaus and Pence (1998) on the changing wealth distribution.

constraint would impose a co-integrating relationship. This is not the case in cross-sectional data, a point made famously by Milton Friedman. Even looking at predicted income, the coefficient is far from unity, suggesting only a thirty percent increase in consumption with income.[46]

The increase in wealth to income explains, again in a partial equilibrium sense, about a fifth of the increase the ratio of consumption to income over the period. The implied increase in consumption due to the changes in wealth to income is calculated as follows. The consumption to income ratio that actually occurred is compared to the consumption to income ratio calculated from the estimated parameters and an unchanging distribution of wealth to income.[47] By estimating the consumption function rather than looking for evidence in Euler equations or contemporaneous relationships, this analysis exploits the long-term relationships between the variables. Thus it finds a significant impact of stock market activity on consumption, where many studies before, focussing on high-frequency data, have found little relationship.[48]

In addition to a role for wealth, the regressions in Table 3 find a significant role for both time and birth-year effects. First, the majority of the increase in the ratio of consumption to income is due to time effects.[49] This is consistent with the "optimism" explanation for the consumption boom in which households believe that future output less government consumption will rise significantly. However, across specification, the null hypothesis that birth year does not belong in the regression model is rejected at the ten percent level across all specifications. Thus, the wealth variable is not sufficient to capture all the cohort effects that are present in the data. The large share of the decline in saving that cannot be explained by the wealth distribution is instead explained by some combination of time effects and unmodelled cohort effects. We can conclude that neither the increase in wealth nor the changing distribution of the population can account for the consumption boom.

The third and fourth regressions investigate the role of stock market participation. If some households are exogenously barred from investing in the stock market, then the consumption of households that are in the market should be higher than those that are outside the market, given the value of the set of state variables for that household.[50] This might be the case if, for example, poor households do not find it worthwhile to pay a fixed costs that is required for access to the stock mar-

[46] It is most likely that this signals persistence but not permanence in the expected/permanent component of income.

[47] The change in the log of wealth less the change in the log of income is multiplied by the estimated coefficient on the log of wealth and added to the change in the fraction of low wealth households times the coefficient on low wealth.

[48] See Poterba and Samwick (1995), Ludvigson and Steindel (1999), and the cites therein.

[49] Changes in the age distribution contribute a small decrease in the consumption to income ratio.

[50] The household that is not excluded can always mimic the excluded household and do at least as well.

ket.[51] Table 3 estimates that the benefits to participation are quite small, on the order of $3-4$ percent of consumption. Given that the share of households in the stock market has risen by about ten percent over the period studied, a partial equilibrium model would predict a third of a percent rise in consumption from increased stock market participation. Of course, in general equilibrium, prices respond. The increased participation affects asset prices and so those already in the market; the expectation of entering the market has effects on those not in the market; and in addition, endogenous changes in the capital stock impact all workers. From this analysis, one can only conclude that there are small but significant increases in consumption from stock market participation above and beyond wealth holding, income, age, and the aggregate state.

In sum, this section finds a significant but small role for the appreciation of assets in the consumption boom: the increase in wealth that occurred from 1984 to 1994 increased the consumption ratio by one fifth of its increase. The remaining causes of the consumption boom are due to other time and cohort effects, but not due to the changing age distribution of the population.

The next section studies the growth rate of consumption and models all time effects as due to the real interest rate or shocks to wealth.

8. Consumption Growth: Impatient Generations, Wealth Increases, and Intertemporal Substitution

This section analyzes the growth rate of consumption instead of its level. The advantages of this are threefold. First, the real interest rate and thus intertemporal substitution is modelled structurally. Second, the growth rate of consumption is related to wealth measures to evaluate whether unexpectedly high asset returns are the cause of the consumption boom. If a series of unexpectedly high stock market returns have increased consumption significantly, the households who own stocks should have significantly higher consumption growth than those that do not. Third, the role of some preference heterogeneity is modelled by allowing different cohorts to have different discount rates.[52] To preview the findings, there is no evidence uncovered that wealthy households had faster consumption growth or that younger cohorts have higher discount rates.

Analysis of growth rates cannot replace examination of consumption levels for two reasons. First, growth rates of consumption at the household level are extremely variable which weakens statistical inference. Second, household transitions like divorce, marriage, death and leaving home imply that the analysis misses significant parts of consumption growth. For example, if young cohorts start life with high consumption and then have consumption growth over their lives that is

[51] See Vissing-Jorgensen (1998).

[52] In the levels analysis, if discount rates were heterogeneous, then the age profiles of consumption would vary with cohort and this variation would undermine the identification employed in section 6.

similar to that of older cohorts, consumption growth aggregated from household consumption growth will show no consumption boom or cohort heterogeneity. The level and the growth-rate analyses are complementary.

Before presenting the analysis, it is important to note that there is a consumption boom in the first-differenced data.[53] However, for the analysis of consumption growth rates, a modified method is used to impute consumption in the $PSID$, as described in the appendix. This imputation assigns $NIPA$ consumption so that the aggregated household data match $NIPA$ growth in real per capita consumption. The imputation does not alter the cross-sectional pattern of consumption growth so that, for example, if stock holders have faster consumption growth than non-stockholders over the period, this will still be detected. This imputation mainly smooths out the swings in growth that occur from year to year due to sampling and measurement error.

The expected real interest rate is constructed from the after-tax nominal return on a six-year Treasury bill during the calendar year of the interview less the inflation rate calculated from the chained deflator for nondurable goods and services that is used to deflate the rest of the data. The marginal tax rate is taken from Stephenson (1998) (the series AMEITRPI).[54] The expectation is taken by predicting the real interest rate for year t (to be used as the return between t and $t+1$) using the following variables: the once lagged second-quarter to second-quarter growth rate in national income; the twice lagged after-tax real interest rate; the once and twice lagged annual unemployment rate for white males 20 years of age and older.[55] The predicting equation is run for the period 1962 to 1997.

Finally, two steps are taken to minimize the impact of the high level of noise in consumption growth data.[56] First, the groupings of age and cohort are expanded to ten-year groups. The noise in consumption growth makes the identification of age and cohort groups more difficult and the five-year groups were substantially noisier.[57] Second, changes greater than or less than 75% in absolute value are dropped.

Identification is slightly simpler in the growth rate regressions. In theory, the innovations in the Euler equation are mean zero and are not predictable by the other right-hand side variables. In other words, the real interest rate captures all time effects that are not orthogonal to cohort and age effects and the real interest rate.

[53] See appendix Figure $A.1$ and Figure 6b.

[54] Using the real return on high grade municipals which are tax-free leads to the same conclusions throughout since the expected returns of these annual series are highly correlated.

[55] The consumption data in the $PSID$ refer to a specific point in time, and are not averages over a calendar year, although there is some debate on this point (see the appendix of Zeldes (1989)).

[56] The same set of regressions are run in grouped data since the measurement error is reduced by averaging, but exogenous variation is also averaged and the results are quite similar to those presented here.

[57] Put another way, the groupings are informally imposing a smoothness prior on the data. Large amounts of variation across neighboring groups suggests insufficient smoothing.

However, one of the main explanations of the consumption boom is that there has been a sequence of positive shocks to wealth. Thus, as a second assumption, time effects aside from the real interest rate are allowed to differ by household wealth holding patterns. That is, the weakness of the first assumption is that innovations to wealth might be correlated with predictable movements in the real interest rate in a short panel of data. Suppose that the period from 1984 to 1994 experienced a run of innovations to wealth, due to unexpectedly strong stock market growth. There would be increases in consumption over the period that would not be mean zero after removing the substitution effect due to movements in the real interest rate. The coefficients on the remaining regressors would suffer from a small-sample bias. To allow for this possibility, I identify the trend in cohort and age effects of all households using the non-stockholders or low-wealth households according to the first identifying assumption and then allow the time effects or tend consumption growth rate of stockholders or high wealth households to be different. This is done by adding a dummy variable for stockholding or the log of wealth to the Euler equation to capture the mean of the expectation errors for these households in sample.

Table 4 shows the results of estimating the following consumption Euler equation[58]

$$\Delta \ln C_{h,t+1} = \sigma E_t[r_{t+1}] + age_{h,t} + cohort_{h,t} + \eta_{h,t+1}$$

where σ is the intertemporal elasticity of substitution. The regressions explain just over one percent of the variation of household consumption growth. The first column of Table 4 presents the regression results for a standard Euler equation.

The first result of interest is that the intertemporal elasticity of substitution is estimated as 0.7. Typical estimates in the literature are significantly lower and sometimes zero.[59] This estimate is in line with Attanasio and Weber (1995) who use grouped CEX data to study Euler equations over the same period. The reasons for this finding here are three. First, consumption of nondurable goods and services typically has a higher elasticity than food. Second, the data are annual. If seasonal fluctuations in consumption and the real interest rate are to some extent driven by preferences, this confounds inference. Finally, for the decade covered by the household data, consumption growth and the expected real interest rate are highly correlated.

Figure 9 displays the expected real interest rate and the growth rate of real consumption per capita. Over the past twenty years, the changes in the growth rate of consumption can be rationalized by movements in the expected real interest rate assuming an intertemporal elasticity of substitution near unity. As to explaining the consumption boom, one can ask to what extent consumption growth would have been slower had a lower real interest rate been in effect.[60] The expected real interest rate from 1980 to 1994 averaged 1.5 percent. During the last five years the

[58] Estimation employs two-stage least squares and reported standard errors account for correlation across households within a time period by including time-effects.

[59] See the discussion in Deaton (1992).

[60] There would of course an associated jump in consumption with an announced different path

expected rate has averaged just over one percent. Given the estimated elasticity of intertemporal substitution, consumption growth would have been 0.35 percent per year slower had this lower interest rate been in effect. Over the 15 years of data on which the coefficient is estimated, consumption grew 5.5 percentage points more than income and this alternative scenario generates nearly exactly that excess.

There are three main problems with explaining the decline in saving solely by intertemporal substitution. First, the nice fit of the Euler equation, roughly since Hall (1978) pointed the equation out, is not evident in the earlier data.[61] Expected income growth may be partly generating this high estimate of the intertemporal elasticity of substitution during the consumption boom. During the 1980's and 1990's, there is a strong correlation between expected income growth and the expected real interest rate.[62] Second, from 1960 to 1979, the real rate of return averaged 0.02 percent, and as shown in Table 1, the growth rate of real consumption per capita averaged 2.5 percent. That is, across the decades, high real interest rates are correlated with low rates of consumption growth. Finally, it is difficult to take seriously a story in which almost none of the movements of consumption over 14 years are driven by changes that represent new information to households.

Despite this scepticism, it is important to note that the consumption and real interest rate data are consistent with the impulse response of a shock to household propensity to consume in the early 1980's.

Turning to the hypothesis that different cohorts have different discount rates, table 4 demonstrates that the cohort effects on consumption growth are small and not significantly different from one another. While the standard errors are large, even in the point estimates, there is not evidence of greater impatience in younger cohorts. It is worth noting that the mean of the cohort dummies is not separately identified from the mean of age effects. Thus one cannot construct a hypothetical consumption path along the lines of Figure 8 without some further restrictions on the data.

The second and third columns of Table 4 show that consumption growth is not significantly higher for high-wealth households, home owners, stockholders, or households with pensions.[63] Wealth is statistically significant in the last column, but the magnitude of the effect is small, suggesting a one-tenth of a percent lower

of interest rates, so this counterfactual is asking whether the observed consumption growth can be rationalized by only the substitution effect.

[61] The usual cites are Hansen and Singleton (1983), Hall (1988), Campbell and Mankiw (1989), and Blinder and Deaton (1985).

[62] Janice Eberly and John Campbell both suggested that I include expected income growth in the consumption growth regressions. Doing so does give a statistically significant role for expected income growth, but it is economically small and does not alter the coefficient on the expected real interest rate. Given the imputations made, this is not quite a fair test of the role of expcted income, but there are many in the literature.

[63] In regressions using wealth data that are only available in 1984, 1989, and 1994, the most recent predetermined value is used. When this is not available, 1984 data are used. Dropping all changes prior to $1984-85$ leads to the same conclusions.

rate of consumption growth for a doubling of wealth.[64]

One possible reason for the insignificant results in these growth-rate regressions is the presence of large amounts of measurement error in the growth rate of consumption. One solution, that comes at the cost of a representative sample, is to regress the growth of consumption in the five years following a wealth survey on the initial wealth levels and time effects and household characteristics such as family size and age, as is done for levels in the previous section. Doing this confirms two of the three main implications of the growth rate regressions. First, cohorts cannot be ignored even after conditioning on the wealth characteristics of households, although it is still not possible to identify a clear pattern of differing discount rates across cohorts. Second, the wealthy are again found to have slightly lower consumption growth over this period. The final main point, that cannot be meaningfully confirmed with only two observations on consumption growth, is that consumption growth and the real interest rate move lock-step.

In sum, how does the analysis of growth rates inform what was learned in the levels analysis? The real interest rate may have played a role, but only as it propagates a positive shock to the desire to consume out of output in the early 1980's. We still find no evidence that the consumption boom is due to wealth appreciation.

9. Conclusion

This paper describes a striking increase in the share of U.S. output that is consumed. This increase occurs concurrently with a reduction in the growth rate of consumption per capita, a high real interest rate, and an increasing ratio of wealth to income. In a search for clues, the paper uses a dataset of household consumption, income and wealth to decompose the consumption boom and confirm or reject possible culprits.

This analysis leads to several conclusions about the large increase in the consumption share of output and the decline in the U.S. saving rate.

First, a thorough examination of $NIPA$ data shows that households and governments in the U.S. are consuming a greater share of output than twenty years ago. Second, this increase is not due to the changing age distribution of the U.S. population.

Third, only one-fifth of the increase in consumption to income can be explained by changes in the ratio of household wealth to income. While the wealth to income ratio has risen, it has done so primarily after the increase in the consumption share of output. The national saving rate has actually risen coincident with the stock market boom of the late 1990's. The propensity to consume out of wealth estimated from the household data cannot rationalize the consumption boom. The increases

[64] While not consistent with the wealthy having more positive innovations to the marginal utility of wealth over this period, it is consistent with the wealthy having lower precautionary saving motives.

in consumption to income ratios across groups are not related to the distribution of wealth, home ownership, or pension participation. While surely they have a role, shocks to asset values are not the main force driving the relative increase in consumption.

Fourth, prime candidates for explaining the consumption boom are factors that increase the effective discount rate of the representative agent. During this period of rising consumption share, the growth rate of real consumption per capita has fallen. At the same time, real interest rates have been relatively high. These two facts together imply a driving force that has increased actual or effective discount rates. It is also worth noting that there is a strong correlation between the real interest rate and consumption growth within the period of consumption boom. That is, the aggregate consumption Euler equation provides a better description of the data during this period than in previous periods.

This paper considers several explanations that can generate this effective impatience. The analysis reveals no evidence that the growth rate of consumption, and thus the discount rate, is higher for younger households. Further, inconsistent with an explanation that relies only on intergenerational government transfers, younger cohorts have a higher rate of consumption to income than older cohorts. Finally, relaxed liquidity constraints could lead to an increase in debt and consumption. But the total increase in debt to income over the past two decades only amounts to one third of the value of the consumption boom.

While we do not yet have a clear answer as to what has caused the recent decline in saving, some speculation is possible based on the concrete findings of this paper.

Given that consumption is a forward-looking variable, households may be learning about high levels of output in the future. This explanation is untestable, and twenty years is a long consumption boom without yet seeing a shift to higher output growth. However, given that other explanations have come up short, this possibility gains credence. The strength of this explanation is that we do observe some signals of high future growth rates, such as the increase in stock prices; the weakness is that without quite a run of negative expectational errors, this explanation cannot match the slowdown in consumption growth.

A second candidate is that rather than being driven by technology or a force external to U.S. households, the decline in saving is due to a shift in the preferences of the typical household. This explanation is as hard to evaluate as the optimism explanation just discussed, however it can fit the facts uncovered here.

A final explanation consistent with the findings of this paper is a combination of factors that work to increase the consumption of different generations. Perhaps federal transfers in the form of Social Security and Medicare are increasing the consumption of the elderly, while relaxed liquidity constraints are allowing the young to consume more of their incomes. This explanation can match the cross-cohort consumption to income effects found in section 6, the high real interest rate and the slowdown in consumption growth; however, it is inconsistent with the stock market boom.

There are many theories that can explain an increase in the consumption of aggregate output. This paper shows that the main monocausal explanations fail to match the household behavior or macroeconomic outcomes observed during the decline in U.S. saving over the past two decades. More importantly, we have an increasing number of facts that new theories or combinations of theories must fit.

References

Alessie, Rob, Arie Kapteyn, and Annamaria Lusardi, "Explaining the Wealth Holdings of Different Cohorts: Productivity Growth and Social Security," Nov 1998. Mimeo.

Attanasio, Orazio, "Personal Saving in the US," in James M. Poterba, ed., *International Comparisons of Household Saving*, University of Chicago Press, 1994.

__ , "Cohort Analysis of Saving Behavior by U.S. Households," *Journal of Human Resources*, 1997, *33* (3), 575–609.

__ , "Consumption and Saving Behavior: Modelling Recent Trends," *Fiscal Studies*, 1997, *18* (1), 23–47.

__ **and Guglielmo Weber**, "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, Dec 1995, *103* (6), 1121–57.

__ **and Martin Browning**, "Consumption over the Life Cycle and over the Business Cycle," *American Economic Review*, Dec 1995, *85*, 1118–1137.

Ayagari, Rao, "Uninsured Idiosyncratic Risk and Aggregate Savings," *Quarterly Journal of Economics*, 1993, *109*, 659–684.

Bernheim, B. Douglas and John Karl Scholz, "Private Saving and Public Policy," in James M. Poterba, ed., *Tax Policy and the Economy*, Vol. 7, MIT Press, 1993.

Bernheim, B.D. and John B. Shoven, eds, Chicago University Press, 1991.

Blinder, Alan S. and Angus S. Deaton, "The Time-Series Consumption Revisited," *Brookings Papers on Economic Activity*, 1985, pp. 465–521.

Bosworth, Barry, "Comment," *Brookings Papers on Economic Activity*, 1996, (1), 391–5.

__ **, Gary Burtless, and John Sabelhaus**, "The Decline in Saving: Evidence from Household Surveys," *Brookings Papers on Economic Activity*, 1991, (1), 183–91.

Branch, E. Raphael, "The Consumer Expenditure Survey: a comparative analysis," *Monthly Labor Review*, Dec 1994, pp. 47–55.

Browning, Martin and Annamaria Lusardi, "Household Saving: Micro Theories and Macro Facts," *Journal of Economic Literature*, Dec 1996, *34* (4), 1797–1855.

Bureau of Economic Analysis, "GNP: An Overview of Source Data and Estimating Methods," Sep 1987. U.S. Department of Commerce Methodology Paper BEA-MP-4, Government Printing Office: Washington, DC.

__ , "Personal Consumption Expenditures," Jun 1990. U.S. Department of Commerce Methodology Paper BEA-MP-6, Government Printing Office: Washington, DC.

__ , "Updated Summary NIPA Methodologies," *Survey of Current Business*, Sep 1997, pp. 12–33.

Bureau of Labor Statistics, *Consumer Expenditure Survey, 1993:*

Interview Survey, Washington, DC: US Department of Commerce, Bureau of the Census, 1980-1993. Ann Arbor, MI: Inter-university Consortium of Political and Social Research.

Caballero, Ricardo J., "Earnings Uncertainty and Aggregate Wealth Accumulation," *American Economic Review*, Sep 1991, *81* (4), 859.

Campbell, John Y. and N. Gregory Mankiw, "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence," in Oliver Jean Blanchard and Stanley Fischer, eds., *N.B.E.R. Macroeconomics Annual*, Cambridge: MIT Press 1989, pp. 185–215.

Carroll, Christopher D., "How Does Future Income Affect Current Consumption," *Quarterly Journal of Economics*, 1994, *109* (1), 111–147.

__ , "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, Feb 1997, *107* (1), 1–56.

__ **and David N. Weil**, "Saving and Growth: A Reinterpretation," in "Carnegie-Rochester Conference Series on Public Policy," Vol. 40 Elsevier Jun 1994.

__ **and Lawrence H. Summers**, "Consumption Growth Parallels Income Growth: Some New Evidence," in B.D. Bernheim and John B. Shoven, eds., *National Savings and Economic Performance*, Chicago University Press 1991.

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

__ , "Saving and Growth," Nov 1997. Manuscript, Research Program in Development Studies, Princeton University.

__ **and Christina Paxson**, "Saving, Growth, and Aging in Tiawan," in David A. Wise, ed., *Studies in the Economics of Aging*, Chicago: Chicago University Press for NBER, 1994, pp. 331–57.

__ **and** __ , "Growth, Demographic Structure, and National Saving in Tiawan," 1997. Manuscript, Research Program in Development Studies, Princeton University.

Denison, Edward, "A Note on Private Saving," *Review of Economics and Statistics*, Aug 1958, *15*, 261–7.

Gale, William G. and John Sabelhaus, "Perspectives on the Household Saving Rate," *Brookings Papers on Economic Activity*, forthcoming.

Gokhale, Jagadeesh, Lawrence J. Kotlikoff, and John Sabelhaus, "Understanding the Postwar Decline in U.S. Saving: A Cohort Analysis," *Brookings Papers on Economic Activity*, 1996, (1), 315–407.

Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption Over the Lifecycle," Dec 1997. Social Science Research Insitute Working paper 9722.

Greenwood, Jeremy and Mehmet Yorukoglu, "1974," May 1996. Manuscript, University of Rochester.

__ **, Zvi Hercowitz, and Per Krusell**, "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, Jun 1997, *87* (3), 342–62.

Gustman, Alan L. and Thomas L. Steinmeier, "Effects of Pensions on Saving: Analysis with Data from the Health and Retirement Study," Aug 1998.

NBER Working paper 6681.

Hall, Robert E., "The Stochastic Implications of the Life Cycle Permanent Income Hypothesis," *Journal of Political Economy*, 1978, *86*, 971–87.

__ , "Intertemporal Substitution in Consumption," *Journal of Political Economy*, Apr 1988, *96.*

Hansen, Lars Peter and Kenneth J. Singleton, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy*, 1983, *91*, 249–68.

Hayashi, Fumio, *Understanding Saving: Evidence from the U.S. and Japan*, Cambridge, MA: MIT Press, 1997.

Hendershott, Patric H. and Joe Peek, "Aggregate U.S. Private Saving: Measures and Empirical Tests," in Richard Lipsey and H. S. Tice, eds., *Measures of Saving, Income and Wealth*, University of Chicago Press, 1988, pp. 185–226.

Hurst, Eric, Ming Ching Luoh, and Frank P. Stafford, "The Wealth Dynamics of American Families, 1984-94," *Brookings Papers on Economic Activity*, 1998, (1), 267–337.

Juster, F. Thomas, James P. Smith, and Frank Stafford, "The Measurement and Structure of Household Wealth," *Labour Economics*, forthcoming.

Ludvigson, Sydney and Charles Steindel, "How Important is the Stock Market Effect on Consumption?," Aug 1999. Federal Reserve Bank of New York Research Paper No. 9821.

Lusardi, Annamaria, "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets," *Journal of Business and Economic Statistics*, Jan 1996, *14* (1), 81–90.

Modigliani, Franco and Richard Brumberg, "Utility Analysis and the Consumption Function: an Interpretation of Cross-Section Data," in Kenneth K. Kurihara, ed., *Post-Keynesian Economics*, New Brunswick, N.J. Rutgers University Press 1956, pp. 388–436.

Parker, Jonathan A., "The Consumption Function Re-Estimated," 1998. manuscript, University of Wisconsin.

Paxson, Christina H., "Saving and Growth: Evidence from Micro Data," *European Economic Review*, Feb 1996, *40* (2), 255–88.

Poterba, James M. and Andrew Samwick, "Stock Ownership Patterns, Stock Market Fluctuations, and Consumption," *Brookings Papers on Economic Activity*, 1995, (2), 295–357.

Sabelhaus, John and Karen Pence, "Household Saving in the 90's: Evidence from Cross-Sectional Wealth Surveys," Dec 1998. Manuscript, Congressional Budget Office.

Stephenson, E. Frank, "Average and Marginal Tax Rates Revisited," *Journal of Monetary Economics*, 1998, *41*, 389–409.

Summers, Lawrence H. and Christopher Carroll, "Why is U.S. National Saving So Low?," *Brookings Papers on Economic Activity*, 1987, *2*, 607–35.

Vissing-Jorgensen, Annette, "Limited Stock Market Participation," 1998. MIT Mimeo.

Zeldes, Stephen P., "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, Apr 1989, *97*, 305–346.

</ref_section>

Appendices

A. The Household Data

A.1. The Panel Study of Income Dynamics

The main relevant features of the $PSID$ are described in the body of the paper. Several remaining issues are noted here.

To ensure that the sample is nationally representative, the oversampled Latino subsample is excluded from analysis.

Figure $A.1$ demonstrates that the ratio of total household food consumption to total household income in the $PSID$ matches well the time-series pattern of total food consumption to National Income in the $NIPA$ data. The $PSID$ ratio is persistently lower by about two and a half percent of income. This is because food consumption in the National Accounts includes food purchased by employers and the government, because income in the $PSID$ includes transfers, and because the $PSID$ seems to underestimate total food consumption expenditures by households. This claim is verified by comparing the amounts inferred from the $PSID$ as compared to the CEX.

The $PSID$ total wealth to income ratio matches the net worth to income ratio in the Flow of Funds data well. Both ratios rise significantly from 1984 to 1989 and are roughly the same in 1989 and 1994.

A.2. The Consumer Expenditure Survey

This section evaluates the relevant features of the CEX data. In order to do the imputation procedure, a household's consumption must be allocated to a quarter and to evaluate the quality of the data it must also be allocated to a year. A household's reported consumption expenditures are allocated to the calendar quarter closest to the mid-point of the year covered by interviews. Annual data are constructed for graphing using the average of all quarters in that year.

Figure $A.1$ shows that the food consumption to income ratio in the CEX declines slightly more and has a sightly lower correlation with the $NIPA$ series than the $PSID$ series does. In fact, this large decline in food consumption to income in the CEX is symptomatic of a poor correlation between the total consumption to income ratio in the CEX and that in the $NIPA$. While this difference is in part due to increasing purchases of medical care by the government, it is also due to an increasing difficulty for the BLS in measuring certain categories of household consumption expenditures. It turns out that this does not create an insurmountable difficulty for the analysis. Instead of taking imputed consumption expenditures as the truth, two adjustments are made so as to allocate $NIPA$ consumption and medical care in relation to imputed household consumption. Gokhale et al. (1996) use the CEX in a similar manner to allocate National Accounts consumption across age groups in each year. In addition, since the CEX is used to "scale-up" food consumption in the $PSID$, the ratio of total nondurables and services

consumption to total consumption, rather than the ratio of consumption to income is the relevant series. The ratio of nondurables and services consumption to food consumption in the CEX and $NIPA$ track each other reasonably well, with the exception of changes both between 1980 and 1982 when the CEX improved its survey instrument for consumption, and between 1986 and 1988.

A.3. The Creation of Nondurable and Services Consumption in the PSID

The details of the regressors in the consumption imputation procedures are as follows. The main regression employs a log-log specification with a cubic polynomial of the log of food consumption. Since there are possibly different returns to scale in the household consumption of food and other items, the variables allowed to shift preferences include nine family size dummies for household sizes 1 through 9 or more. The imputation also includes dummies for whether the household head has a high-school degree or less education, some college, or a college degree or more education. To account for shifting preferences across ages, I also include a fifth-order polynomial in age for households less than age 65 and second order polynomial in age for households greater than 65. To account for labor supply interactions the preference shifters include a retirement dummy variable, a dummy variable for whether the household is retired and younger than 65, and dummies for whether there are zero, one or two or more earners. Finally, to capture both prices and preferences, a set of quarter dummies and a set of year dummies are included.

The four steps of the imputation are as follows.

First, using the CEX data, the log of nondurable and services consumption is regressed on a cubic polynomial in the log of food consumption and the remaining regressors just discussed. The CEX regression using $37,730$ households explains 80 percent of the variation in household consumption, although the typical error is 30 percent of nondurable and services consumption. The coefficients are not reported but are reasonable. A household with a college educated head consumes fifteen percent more nondurable and services consumption relative to food than a household with a head without a high-school degree. Retired households consume 10 percent more nondurable and services consumption relative to food than a non-retired household.

Second, the estimated equation is used with the same set of regressors in the $PSID$ to predict nondurable and services consumption for each household. The number of earners in the $PSID$ is calculated from reports on labor income and wages of head and spouse. The quarter dummy is set equal to the second quarter, since most $PSID$ households are interviewed in May. Similarly, the year dummy for 1979 is set equal to its value for 1980 and the year dummy for 1994 is set equal to its value for 1993. Constructing the implied consumption to income ratio from the imputed data gives a highly volatile series. This said, the average ratio from the first four years is 0.057 below the ratio for the last four, showing a reasonably good mapping to the aggregate trend.

Third, the imputed consumption for each household is treated as a relative con-

sumption level and the total consumption across households is scaled up to include medical purchases by the government. Medical care purchased by the government, except for Medicare, is allocated in proportion to total consumption across all households that are younger than 65 by year. This adjusts consumption of these households upwards by from 1.5 to 2.5 percent of total consumption over the entire sample. Medicare expenditures are allocated evenly across all households age 65 or older in a similar manner, which leads to a scale factor that grows by 10 percentage points over the sample. The elderly account on average for 11 percent of total imputed consumption. Medicare purchases by the government rise from 1.6 to 3.2 percent of total consumption expenditures less government spending on health care. Without this adjustment, the consumption of the elderly would be significantly understated and, more importantly, the rise in their consumption would be understated. Income is not adjusted for this consumption that is purchased by the government for households. Interpretation of cohort and age profiles throughout the paper keeps this in mind.

Fourth, the consumption of nondurable goods and services in the $NIPA$ in each year is allocated across households in proportion to each household's consumption from the third step. The allocation is conducted so that the consumption to income ratio in the micro-data matches that in the $NIPA$ in every year.

When working with the growth rate of consumption the following modification to the imputation procedure is made. Instead of using the level of predicted consumption in the $PSID$ to allocate $NIPA$ consumption expenditures, the level is used only to allocate only medical purchases by the government. In the fourth step of the imputation, the growth rate of $NIPA$ real consumption per capita is allocated across households in accord with their household growth rates. One might be concerned because this procedure ignores that fact that these two series might differ due to household births and deaths. However, in the $PSID$ data, many missing consumption growth rates are not due to birth or death but to missing data. Thus it is also not appropriate to assume that the difference between the $PSID$ growth in consumption and that in the $NIPA$ represents differences in true births and deaths. More importantly, the trend in the time-series of consumption growth from the $PSID$ is similar whether one calculates it from averaging levels or averaging first differences. See also the discussion in the text.

Table 1

Consumption Growth and Expenditure Shares

Real Per Capita Annual Growth Rate:	**1959-69**	**1969-79**	**1979-89**	**1989-98***
Total PCE	3.0%	2.4%	2.0%	1.6%
Nondurable Goods and Services	2.8%	2.3%	1.8%	1.4%

Change in Share of GDP:	**1959-69**	**1969-79**	**1979-89**	**1989-98***
Total PCE	-1.2%	0.7%	3.8%	2.1%
Durable Goods	0.3%	-0.4%	0.3%	-0.2%
Nondurable Goods	-3.5%	-1.4%	-3.0%	-1.9%
Services	2.0%	2.5%	6.5%	4.2%
Medical Care	1.3%	1.6%	2.6%	1.7%
Other Services	0.8%	0.3%	2.3%	1.9%
Housing	0.0%	0.0%	1.2%	0.2%

* Note 1998 estimates are preliminary.

Table 2

Cell Sizes for Age and Cohort Groups

Cohort Born	Cell Size	Age Group	Cell Size
05-09	1,722	19-24	7,090
10-14	2,799	25-29	12,838
15-19	3,264	30-34	13,075
20-24	3,973	35-39	10,427
25-29	4,725	40-44	7,319
30-34	4,488	45-49	5,263
35-39	4,135	50-54	4,702
40-44	5,462	55-59	4,421
45-49	9,331	60-64	4,225
50-54	13,024	65-69	3,850
55-59	13,627	70-74	3,063
60-64	8,339	75-85	3,208
Partially Observed Cohorts:		85+	583
1893-05	1,176		
65-73	3,999		

Table 3

Consumption Function Regressions

Regression:	1	2	3	4
Log of Wealth (if not low)	0.048	0.039	0.045	0.036
	(0.002)	(0.002)	(0.002)	(0.002)
Low Wealth	0.330	0.267	0.301	0.248
	(0.025)	(0.023)	(0.025)	(0.024)
Expected Log Income	0.314	0.175	0.309	0.172
	(0.005)	(0.006)	(0.006)	(0.007)
Pension	0.049	0.028	0.047	0.027
	(0.007)	(0.007)	(0.007)	(0.007)
Log Income		0.161		0.160
		(0.005)		(0.005)
Stockholder			0.041	0.028
			(0.008)	(0.007)
Year 1989	0.030	0.030	0.028	0.029
	(0.008)	(0.007)	(0.008)	(0.007)
Year 1994	0.053	0.059	0.048	0.056
	(0.008)	(0.008)	(0.008)	(0.008)
Number of Observations	11,903	11,903	11,901	11,901
R-Squared	0.583	0.623	0.584	0.624
Singificance level for birthyear:	0.953	0.908	0.955	0.909
Implied increase in C/Y due to				
Increase in W/Y: 84 - 89	0.011	0.009	0.010	0.008
89 - 94	0.003	0.002	0.002	0.002
Total Increase in C/Y 84 - 89	0.025	0.025	0.025	0.025
89 - 94	0.030	0.030	0.030	0.030

Regressions also include family size and the number of children in the household and a complete set of age group effects.

Table 4

Consumption Growth Regressions

Regression:	1	2	3
Expected Real Interest Rate	0.700	0.729	0.730
	(0.120)	(0.125)	(0.125)
Cohort <09	-0.006	-0.006	-0.005
	(1.259)	(1.308)	(1.309)
Cohort 10-19	0.006	0.010	0.010
	(1.178)	(1.223)	(1.224)
Cohort 20-29	-0.009	-0.005	-0.005
	(1.065)	(1.104)	(1.105)
Cohort 30-39	-0.013	-0.011	-0.011
	(0.918)	(0.952)	(0.953)
Cohort 40-49	-0.009	-0.009	-0.009
	(0.791)	(0.825)	(0.826)
Cohort 50-59	-0.007	-0.008	-0.008
	(0.577)	(0.599)	(0.600)
Log Wealth/100 (if not low)		-0.076	-0.096
		(0.320)	(0.433)
Stockholder		0.000	0.000
		(0.023)	(0.023)
Pension		0.003	0.003
		(0.032)	(0.032)
Homeowner			0.003
			(0.040)

Dependent variable is the first difference of log consumption.
Regressions also include a complete set of age group effects.

Figure 1a

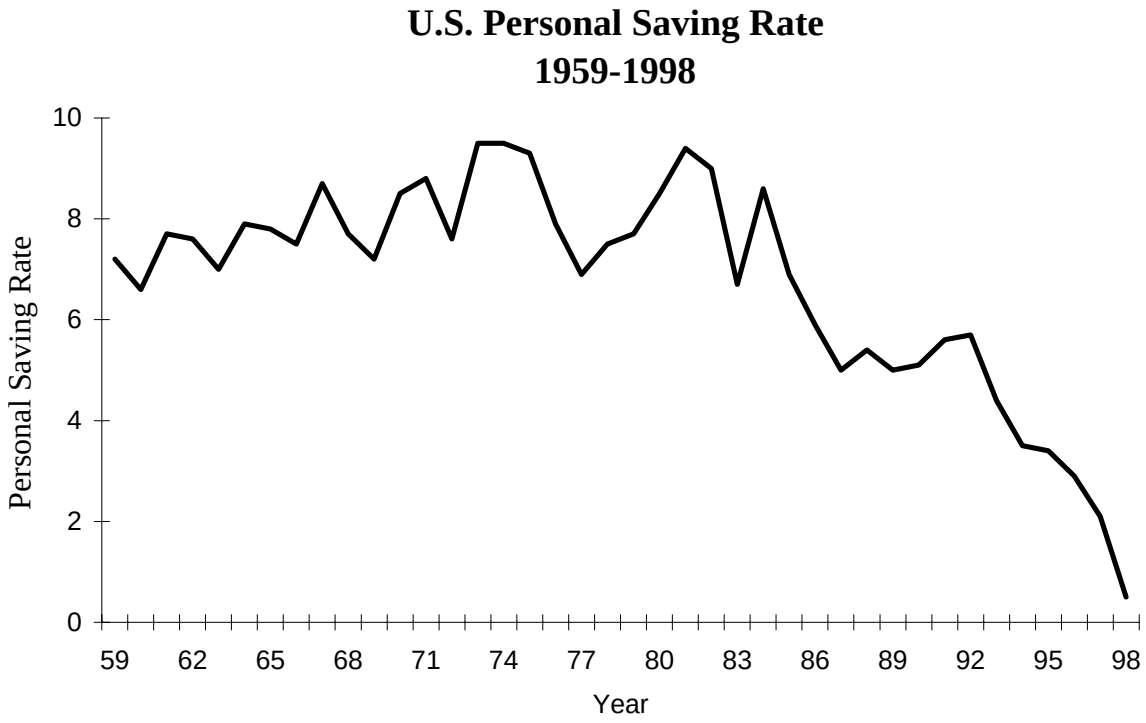

Figure 1b

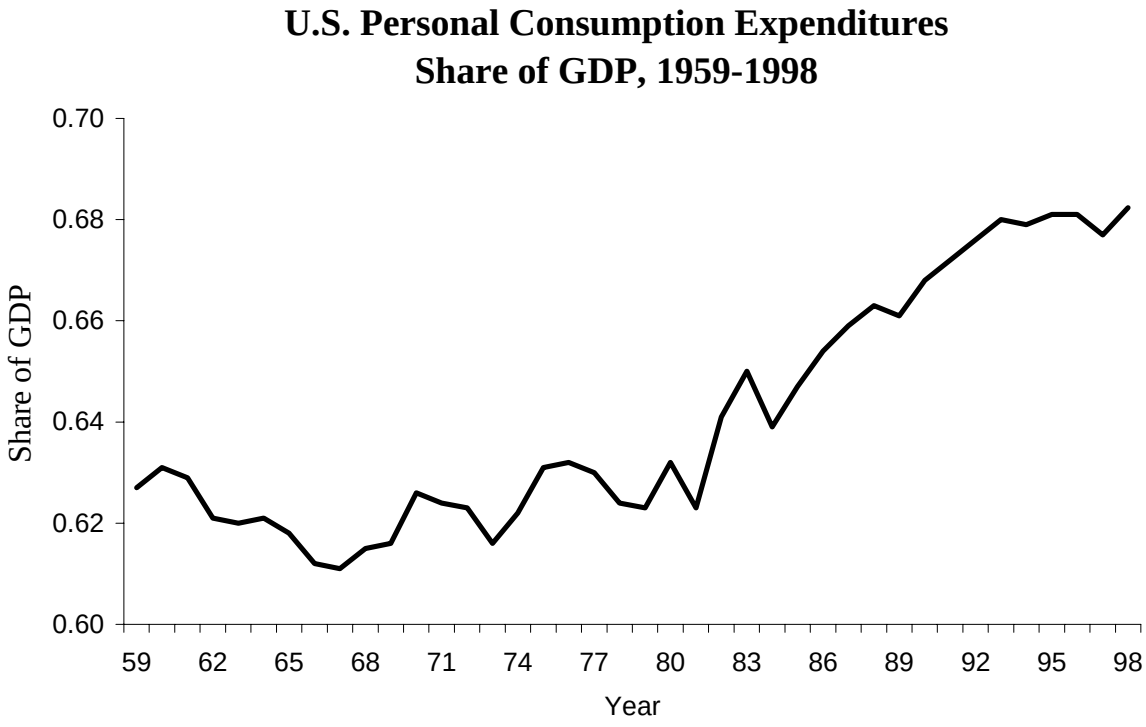

Figure 2a

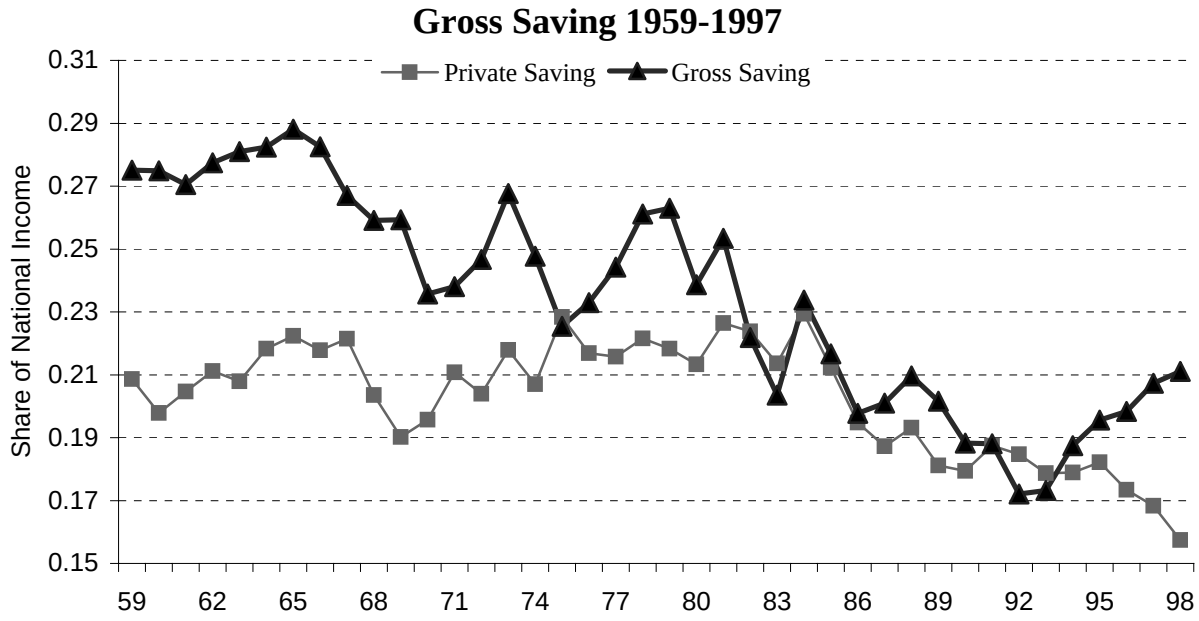

Figure 2b

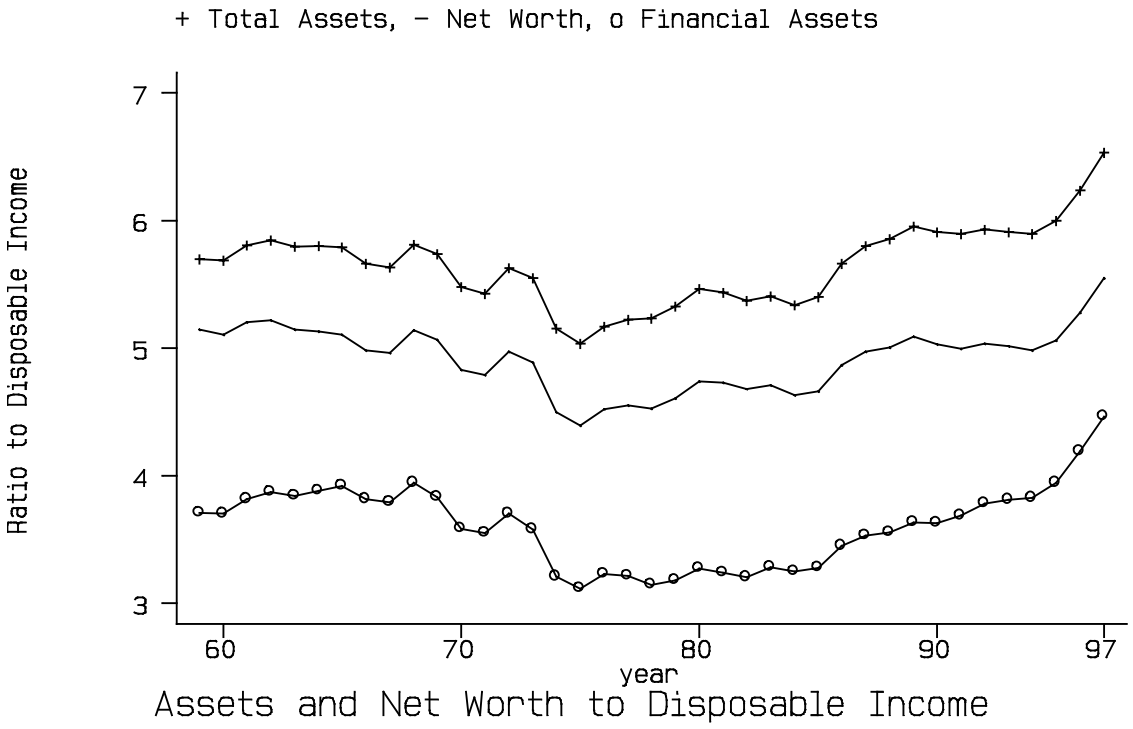

Figure 3

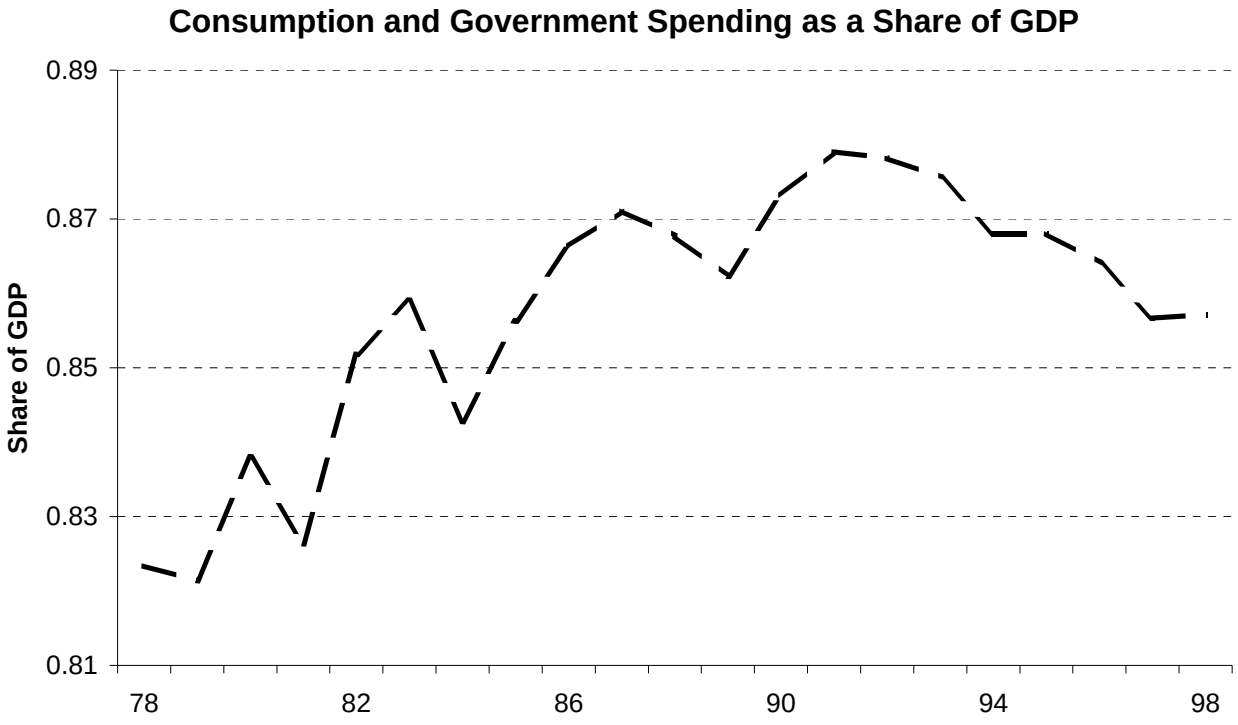
Consumption and Government Spending as a Share of GDP
Share of GDP
0.89
0.87
0.85
0.83
0.81
78
82
86
90
94
98

Figure 4a

Figure 4b

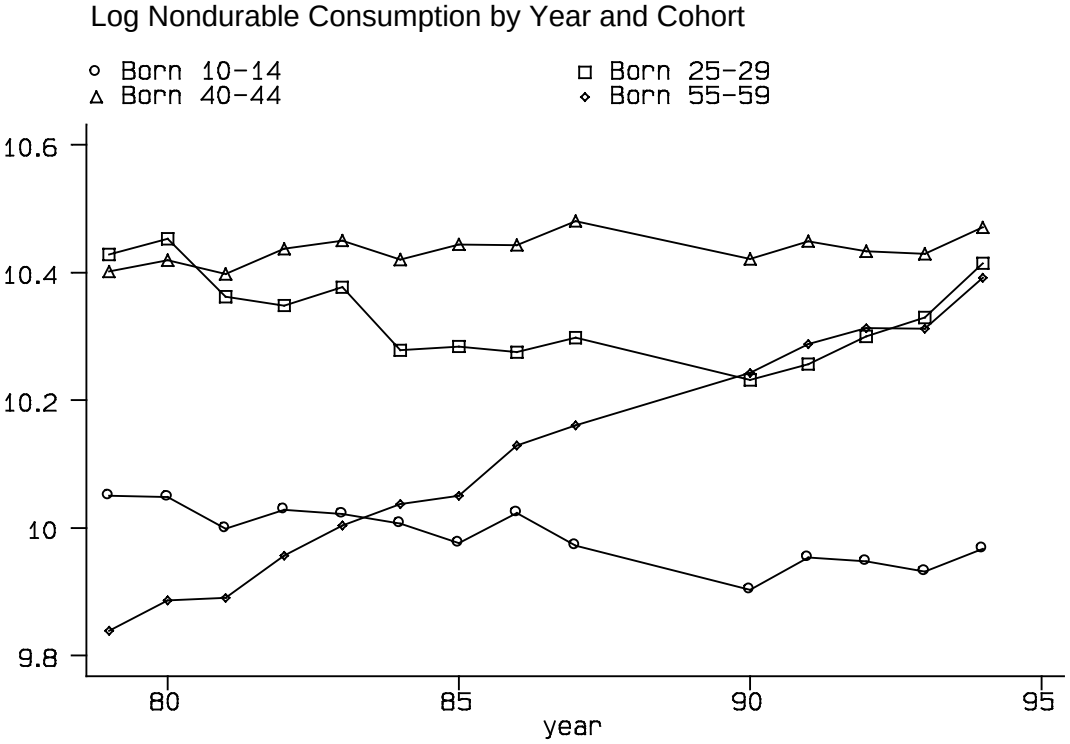

Figure 5a

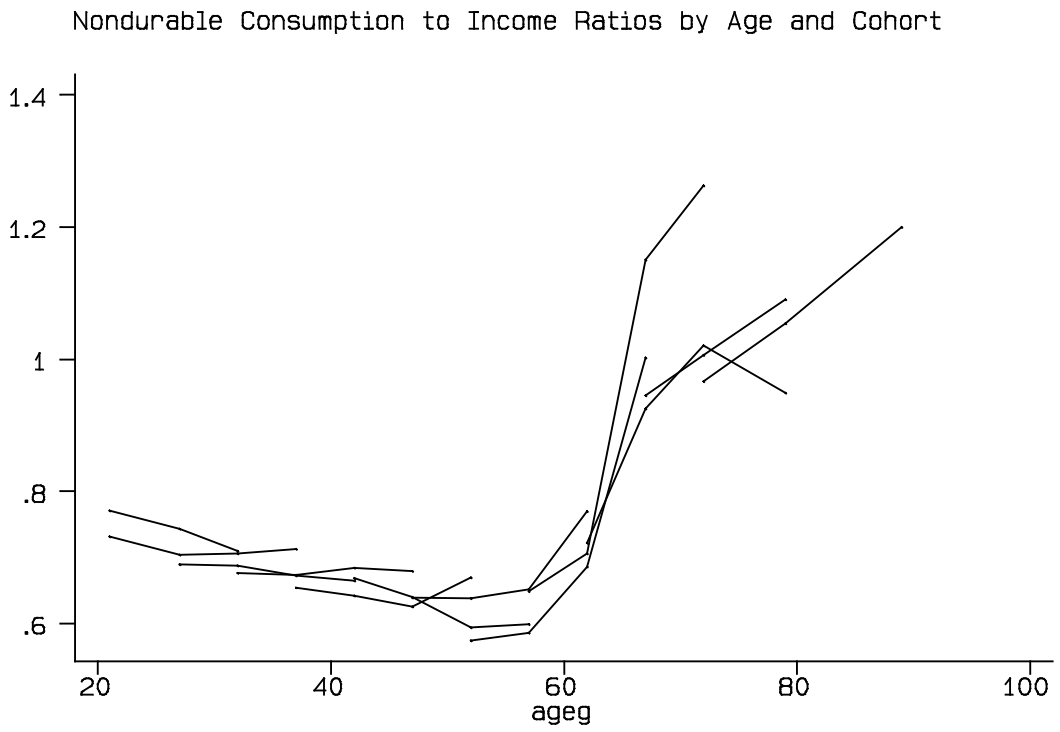

Figure 5b

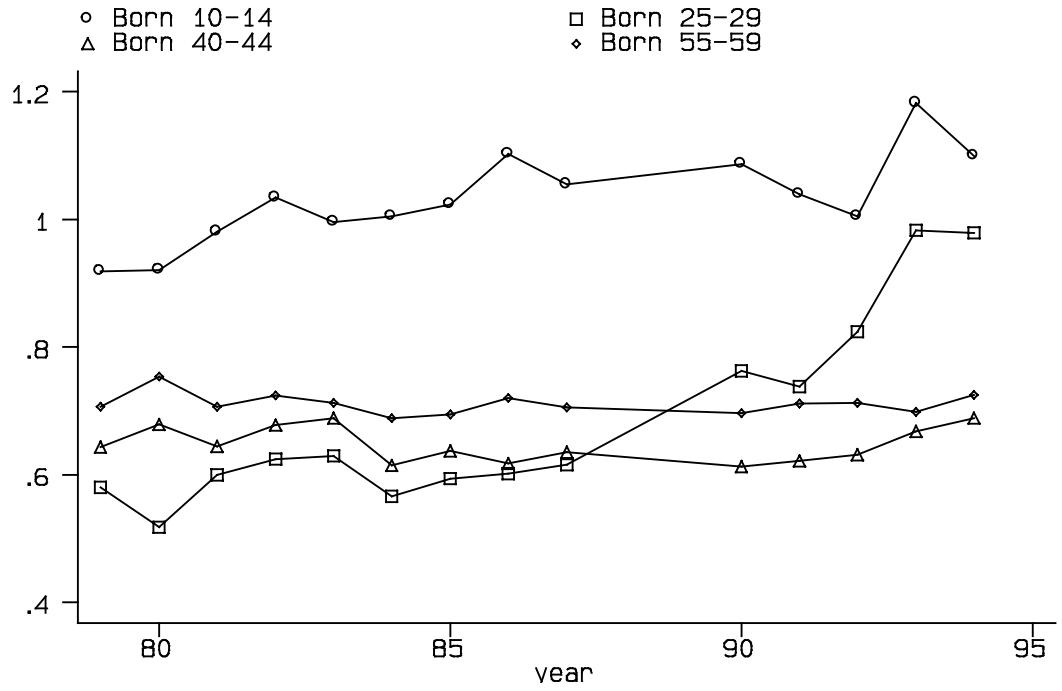

Figure 6a

Figure 6b

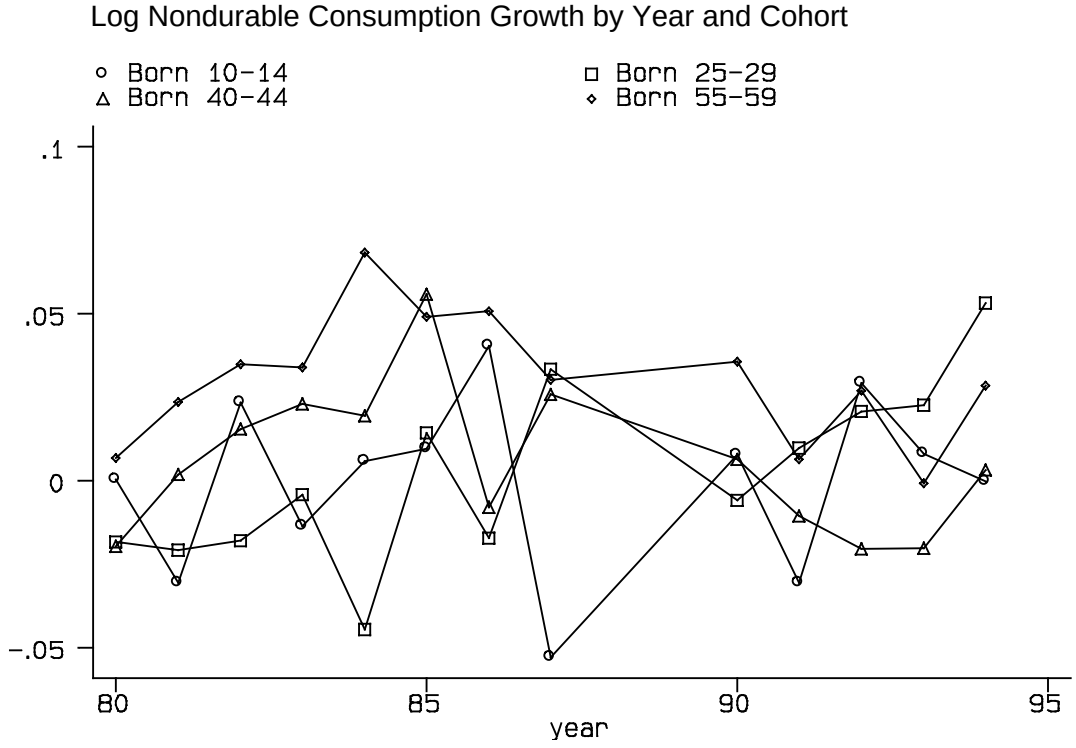

Figure 7a

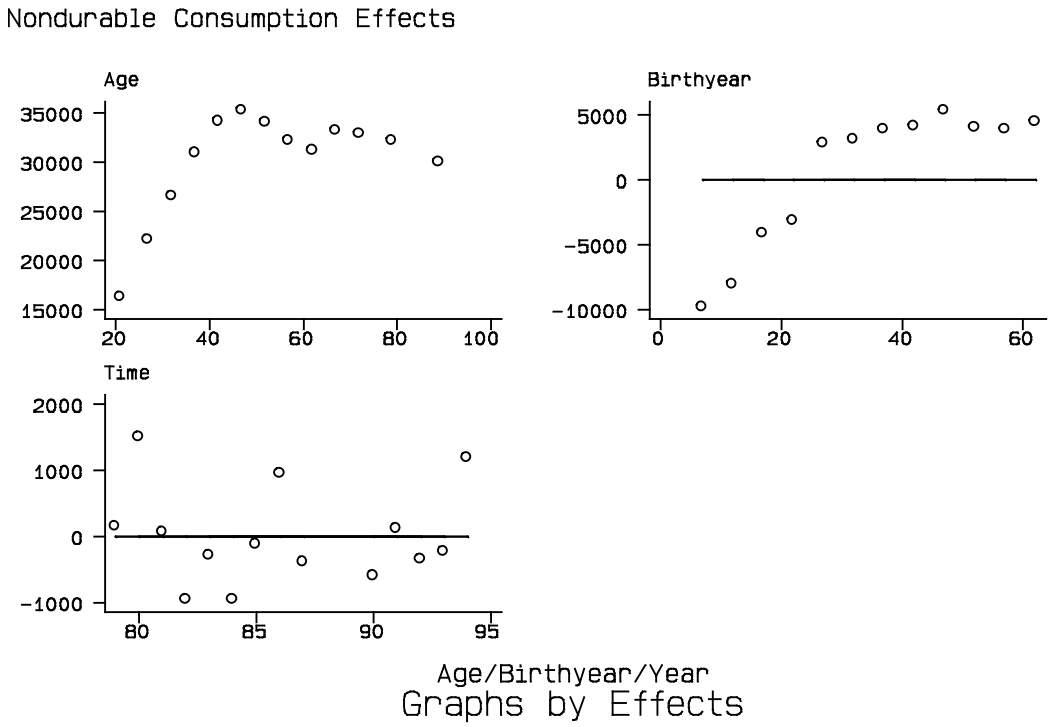

Figure 7b

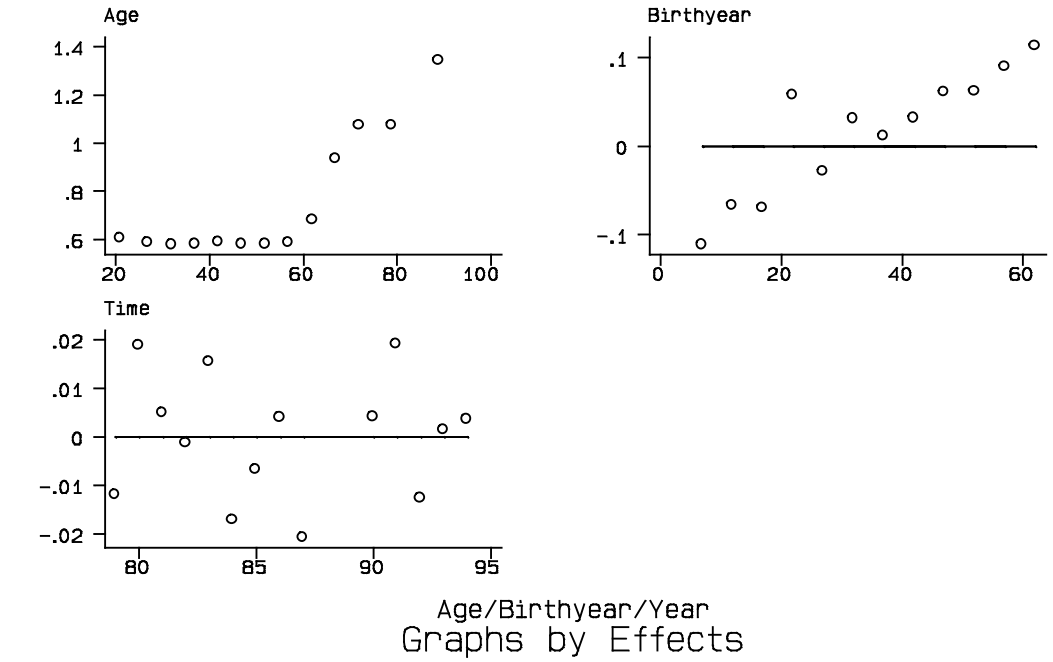

Figure 8a

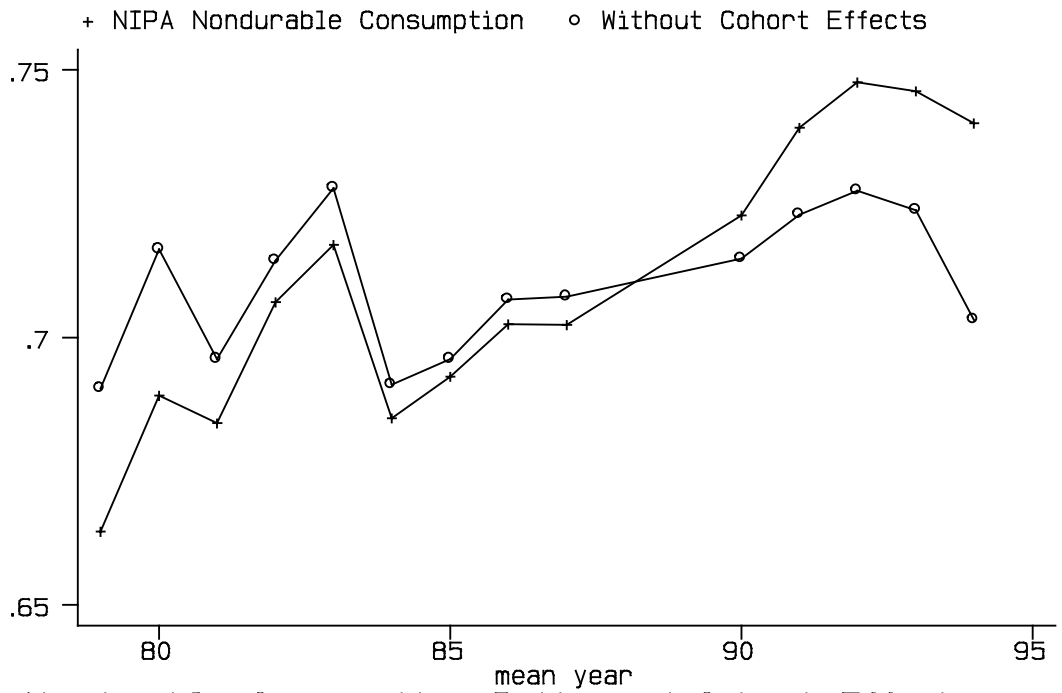

Nondurable Consumption Ratio and Cohort Effects

Figure 8b

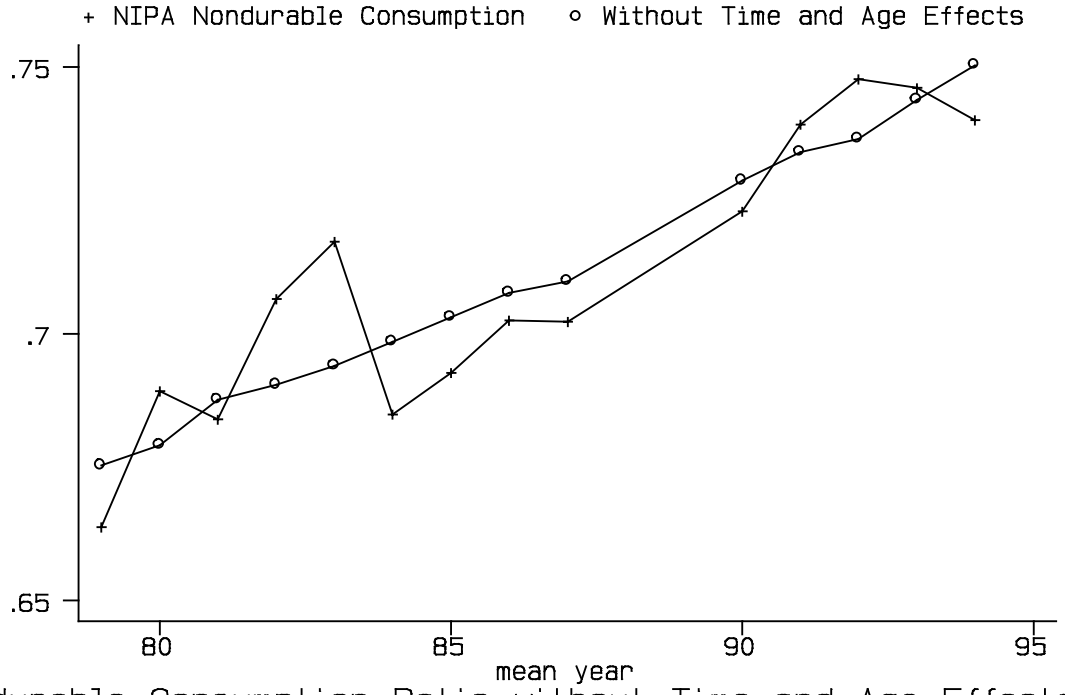

Nondurable Consumption Ratio without Time and Age Effects

Figure 9

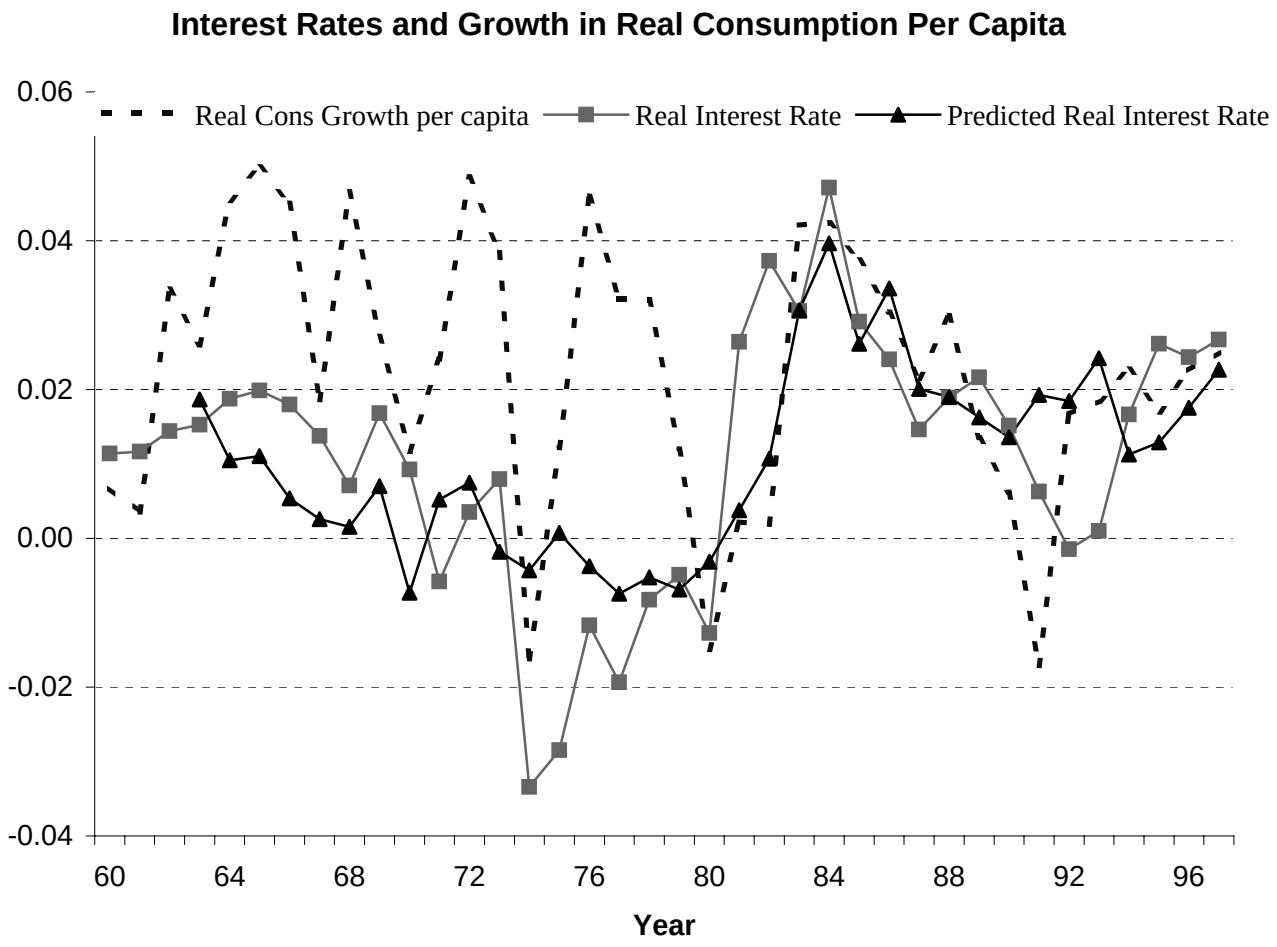
Interest Rates and Growth in Real Consumption Per Capita
Real Cons Growth per capita
Real Interest Rate
Predicted Real Interest Rate
0.06
0.04
0.02
0.00
-0.02
-0.04
60
64
68
72
76
80
84
88
92
96
Year

Figure A.1

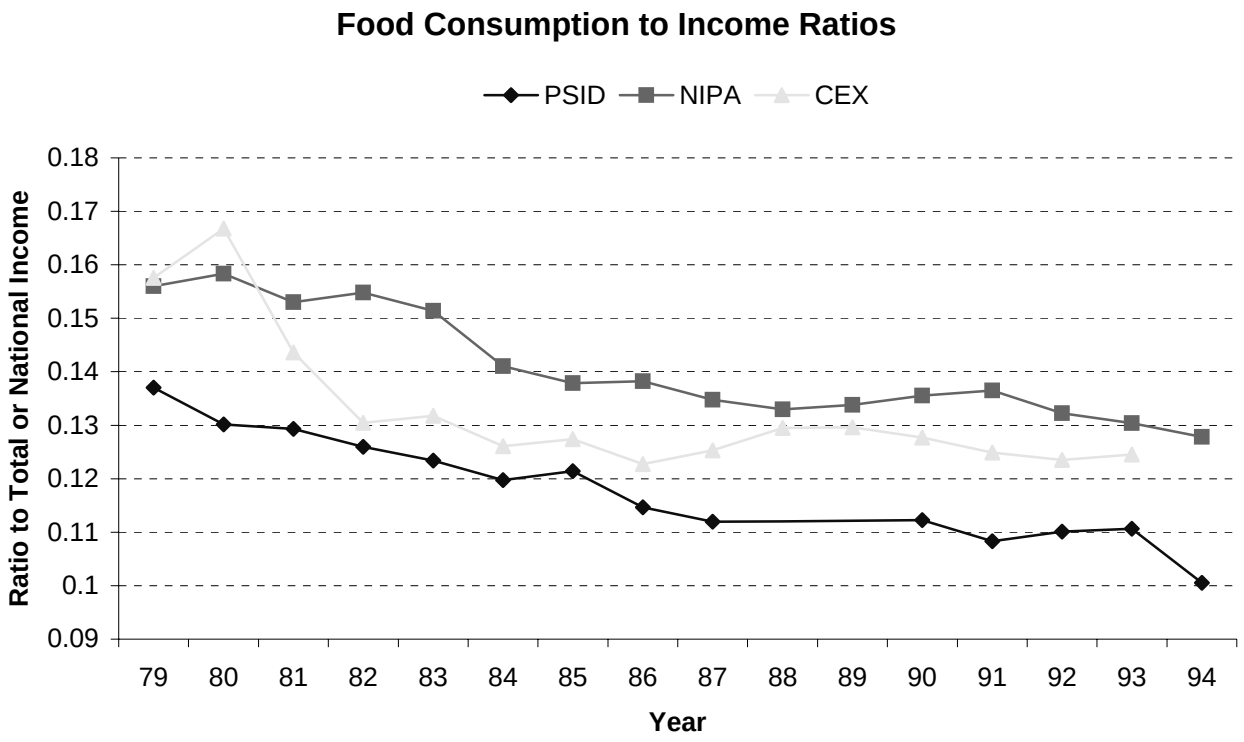
Food Consumption to Income Ratios
PSID
NIPA
CEX
Ratio to Total or National Income
0.18
0.17
0.16
0.15
0.14
0.13
0.12
0.11
0.1
0.09
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
Year