

The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence

Christopher D. Carroll; Robert E. Hall; Stephen P. Zeldes

Brookings Papers on Economic Activity, Volume 1992, Issue 2 (1992), 61-156.

Stable URL:

http://links.jstor.org/sici?sici=0007-2303%281992%291992%3A2%3C61%3ATBTOSS%3E2.0.CO%3B2-N

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/about/terms.html. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

Brookings Papers on Economic Activity is published by The Brookings Institution. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at http://www.jstor.org/journals/brookings.html.

Brookings Papers on Economic Activity ©1992 The Brookings Institution

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor-info@umich.edu.

©2003 JSTOR

http://www.jstor.org/ Wed Jul 23 17:28:35 2003

CHRISTOPHER D. CARROLL Board of Governors of the Federal Reserve System

The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence

As I wRITE, the U.S. economy remains mired in the slowest recovery from any recession in the postwar period. Consumer confidence and consumption spending, in particular, have been exceptionally weak, and the unemployment rate has continued to rise long after many other indicators began to improve. This paper presents evidence that these facts are related, in the sense that consumer pessimism about unemployment explains a substantial part of the recent weakness in consumption. However, neither theoretical consumption models commonly used for macroeconomic research¹ nor standard macroeconometric forecasting models² provide a direct role for unemployment expectations in determining current consumption.³ By contrast, this paper shows that the "buffer-stock" model of saving that has evolved from the work of Stephen Zeldes and Angus Deaton and from my previous work can imply a central role for unemployment expectations.⁴ In the buffer-stock model,

I am grateful to Angus Deaton, Robert Hall, Miles Kimball, Jennifer Manning, Edwin Sanders, Lawrence Summers, David Wilcox, Stephen Zeldes, the members of the Brookings Panel, and participants in seminars at the Federal Reserve Board, the National Bureau of Economic Research, the University of Pennsylvania, and the Bank of Italy. Remaining errors are my own. The views expressed in this paper are my own and do not necessarily reflect the opinions of the Federal Reserve Board.

1. Specifically, I have in mind the certainty-equivalence model used in different ways by Flavin (1981), Campbell (1987), Campbell and Mankiw (1989), and many others.

2. See, for example, the MIT-Penn Social Science Research Center (MPS) model of the consumption sector. See Fuhrer (1992) for a description.

3. In the certainty equivalence model, only the future level of income affects the current level of consumption. Although income and unemployment are likely to be correlated, once expected income is controlled for, the model implies that there is no independent influence of the unemployment rate on consumption.

4. Zeldes (1989a), Deaton (1991), and Carroll (1992).

consumers hold assets mainly so that they can shield their consumption against unpredictable fluctuations in income; unemployment expectations are therefore important because typically the most drastic fluctuations in a household's income are those associated with spells of unemployment.

Buffer-stock saving behavior can emerge from the standard dynamic optimization framework when consumers facing important income uncertainty are both *impatient*, in the sense that if income were certain, they would like to borrow against future income to finance current consumption, and *prudent*, in Miles Kimball's sense that they have a precautionary saving motive.⁵ The buffer-stock behavior arises because impatience makes consumers want to spend down their assets, while prudence makes them reluctant to draw down assets too far. In my 1992 paper, I showed that under plausible circumstances this tension will imply the existence of a target wealth stock.⁶ If wealth is below the target, fear (prudence) will dominate impatience and the consumer will try to save, while if wealth is above the target, impatience will be stronger than fear and consumers will plan to dissave. Unemployment expectations are important in this model because when consumers become more pessimistic about unemployment, their uncertainty about future income increases, so their target buffer-stock increases, and they increase their saving to build up wealth toward the new target.

The model is structurally similar to that of Stephen Zeldes⁷; the main formal difference is the buffer-stock model's assumption that consumers are impatient. Even more similar is Angus Deaton's model, except that Deaton imposes liquidity constraints, while the model in this paper does not.⁸ Indeed, the term "buffer-stock" saving is borrowed from Deaton's description of his own model, and in practical terms I view the two models as relatively close substitutes. However, while Deaton did not treat unemployment in his model, the simulation evi-

5. Kimball (1990a). The theory described in this paper is not directly related to the buffer-stock model of liquidity preference developed, for example, by Cuthbertson and Barlow (1991). Their model is largely an econometric framework for analyzing aggregate consumption data, and is not based on the solution to a consumer's optimization problem. Nonetheless, much of the empirical evidence they provide can be interpreted as supportive of the model presented here.

6. Carroll (1992).

7. Zeldes (1989a).

8. Deaton (1991).

dence presented here suggests that unemployment expectations are probably a crucial factor in determining the amount and characteristics of buffer-stock saving.

The next part of this paper is devoted to analyzing the nature and degree of uncertainty households face, so that the model can be parameterized and solved. Using household data from the University of Michigan's Panel Study of Income Dynamics (PSID), I find a high degree of income uncertainty; in particular, I find that households occasionally experience very bad outcomes in which their income drops essentially to zero.⁹ The third section of the paper solves the buffer-stock model using these uncertainty estimates and explores some of the model's properties. Simulations show that, even with unchanging expectations about the average future level of income, changes in the expected probability of "bad events" (interpreted as unemployment) have a major impact on current consumption and saving. This section also shows that in the buffer-stock model (in contrast to permanent-income models) the elasticity of saving with respect to the growth rate of income is positive. Furthermore, while standard life-cycle and permanent-income models imply that the interest elasticity of saving should be strongly positive,¹⁰ in the buffer-stock model the interest elasticity of saving is approximately zero. Both these results arise because of the target-saving character of the model.

The fourth section provides a variety of macroeconomic evidence that supports the buffer-stock model, as well as some evidence that documents an apparent limitation of the model. The model is shown to be capable of explaining the "consumption/income parallel" that Lawrence Summers and I found in our 1991 paper.¹¹ The model can also explain evidence that shows that consumers express a greater desire to save, and actually save more, in periods when fears of unemployment are high or rising, even controlling for expectations about income growth. The model has somewhat greater difficulty explaining a negative short-term correlation between consumption growth and unemployment expectations, but I argue that the empirical results would be roughly consistent with the model if consumption adjusts slowly, rather than instantaneously, to changes in unemployment expectations.

^{9.} University of Michigan, A Panel Study of Income Dynamics, Wave XVIII.

^{10.} See, for example, Summers (1981).

^{11.} Carroll and Summers (1991).

The fifth section uses the theoretical and empirical results of the previous sections to analyze two puzzles: the secular decline of the personal saving rate in the United States over the last two decades, and the surprising weakness of consumption growth in the past year or two. The last section of the paper indicates directions for future research.

Estimates of Income Uncertainty from the PSID

Unfortunately, most existing microeconomic studies of the structure of changes in income are not useful for characterizing the uncertainty facing consumers in a standard life-cycle/permanent-income hypothesis (LC/PIH) model because they concentrate on inappropriate data samples and inappropriate measures of income.¹² I therefore made new estimates using the PSID, which has data for a large sample of households for the years 1968–85.

The most natural way to link the PSID data to a model of consumption is to identify the PSID "household" as the decision unit and to examine variability in household income. The appropriate measure of income for solving the model below is total family noncapital income, which includes transfer payments such as unemployment insurance, disability compensation, alimony, and social security payments, as well as noncapital income earned by household members other than the head. In all of the analyses below, the sample will consist of households in which the same individual was head of the household over the entire period; the head was older than 24 and younger than 63 years old over the entire sample period; and the household was not part of the PSID poverty subsample.¹³ The sample will also sometimes be restricted to heads of

12. In particular, MaCurdy (1982) and Abowd and Card (1989), two widely cited sources on micro-level income changes, restrict their samples to individuals whose income neither falls to zero nor rises too much, eliminating at a stroke the consumers who experience the greatest income variability. Their focus on labor income of the head of household, on the other hand, exaggerates the variability of total noncapital income because unemployment insurance, disability payments, and labor market participation by other household members can all act to reduce the variation in total family income.

13. Participants in the PSID consist of two groups: a random probability sample of the population as a whole, and a special subsample chosen to study poor people. If the behavior of the poverty subsample differs systematically from that of the rest of the population, then results could be biased by including this subsample in the analysis. Thus I excluded this subsample.

households whose marital status never changed or who were never selfemployed. Restrictions like these will eliminate important sources of income uncertainty, but in general it is best to err on the side of underestimating uncertainty. If true uncertainty is larger than estimated, most of the results of the latter portion of the paper would only be strengthened.

To solve the models below, I will describe the process that generates total family noncapital income YL ("labor income" for short) by the following equation:

(1)
$$YL_t = P_t V_t,$$

where V_t is a multiplicative transitory shock in year t and P_t is "permanent labor income" in year t, the value of labor income if no transitory shocks occur ($V_t = 1$). Note that this definition of "permanent labor income" does not correspond to the common modern interpretation in which permanent income is some function of the present discounted value of all expected future labor income. This definition of P_t does correspond, however, fairly closely to what Milton Friedman himself appeared to mean by permanent income. Friedman writes:

The permanent component [depends on] . . . the nonhuman wealth [the unit] owns; the personal attributes of the earners in the unit, such as their training, ability, personality; the attributes of the economic activity of the earners, such as the occupation followed, the location of the economic activity, and so on. *It is analogous to the "expected" value of a probability distribution.* The transitory component is to be interpreted as reflecting all "other" factors, factors that are likely to be treated by the unit affected as "accidental" or "chance" occurrences.¹⁴

The log of *P* will be assumed to follow a random walk with drift:

(2)
$$P_{t+1} = GP_t N_{t+1}$$
, or $\ln P_{t+1} = \ln G + \ln P_t + \ln N_{t+1}$,

where G is the growth factor for permanent income, and N_t is period t's multiplicative shock to permanent labor income.

The goal of this section is to characterize the distributions of V_i and N_i , the transitory and permanent shocks, respectively. Most of the data will be approximately consistent with the assumption that V_i and N_i are distributed lognormally. The lower tail of the income distribution, how-

^{14.} Friedman (1957, p. 21). Italics added.

ever, is not well captured by the assumption of lognormality, and therefore must be examined separately.

Characterizing the Lower Tail

As a first step toward understanding the structure of the income data, detrended noncapital income YL_{ii} was calculated for each household *i* in each year *t* over the 1976–85 period for which noncapital income can be calculated.¹⁵ For each household, an average level of permanent income \overline{YL}_i was constructed, equal to the average value of YL_{ii} over the 1976– 85 period. Finally, a variable $YLRATIO_{ii} = YL_{ii}/\overline{YL}_i$ was created. This generated ten observations for each household, and sample restrictions narrowed the sample to 1,238 households, yielding a total of 12,380 observations of YLRATIO. Figure 1 presents two histograms. The top figure plots the *YLRATIO* for the entire sample, while the bottom figure magnifies the lower portion of the histogram and illustrates that a substantial concentration of observations occurs at approximately zero income.¹⁶

The observations near zero income do not appear to be drawn from the same distribution generating the rest of the histogram. This encourages the belief that the total distribution can be generated by combining a process that occasionally generates zero income with a well-behaved process that operates whenever income is not zero. Previous researchers have sometimes argued that observations of near-zero income reflect measurement error and hence should be excluded from analysis. Inspection of the raw data, however, shows that, at least for the consumers

15. The detrending procedure was as follows. In each year, mean weighted family income was calculated for the whole sample, and "detrended" household income for each household was defined as the ratio of actual household income to the mean. This procedure removes common movements in income across all households, including both cyclical and secular trend effects. The next step was to remove predictable life-cycle movements from income. An equation was estimated relating detrended income to age, occupation, education, and interactions of these terms. This equation was then used to predict income in each year for each household. Detrended income from the first procedure was then divided by predicted income for each household, generating the *YL* series used in the calculations above. This second procedure should have removed changes in income that were due purely to the predictable effects of aging. Results were basically the same as those reported using a variety of other procedures for removing aggregate and life-cycle trends.

16. There were a few observations of negative noncapital income, but only among selfemployed households. These observations were counted as zero-income observations.



Number of observations



Number of observations





Source: Author's calculations based on data from the University of Michigan, *Panel Survey of Income Dynamics* (PSID), as described in the text. Observations are equal to the number of households multiplied by the number of years in the sample period.

Heads of household	Total number of observations	Number of near-zero events	Percent of near-zero events
Full sample	12,380	81	0.65
No self-employed	10,674	39	0.37
No change in marital status No self-employed or change	11,934	73	0.61
in marital status	10,281	31	0.30

Table	1.	The	Number	of Times	; Family	Noncapital	Income	Falls to	Zero
_				~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~					

Source: Author's calculations using University of Michigan, A Panel Study of Income Dynamics (PSID). See text for definition of various samples.

who are not self-employed, some event such as unemployment, injury, or health problems usually occurred in the same period as, or just prior to, the near-zero event. Furthermore, in a careful study of errors in labor survey data, Greg Duncan and Daniel Hill found that outlying observations on annual income generally correspond to actual extreme experiences of the respondent, not to measurement error.¹⁷

Table 1 counts the number of near-zero income events when the sample is restricted in various ways, where a near-zero event is defined as a value of *YLRATIO* of less than 0.1 (noncapital income less than 10 percent of its average over the period). In the unrestricted sample of 12,380 observations, 81 near-zero events occur. If near-zero income events are independent and transitory and if all consumers are identical, this means that each consumer has about a 0.65 percent chance of experiencing a near-zero event¹⁸ in any given year.¹⁹ This figure is reduced when the sample is restricted to heads of households who were not self-employed or to those who did not experience any change in marital status over the period, but it is not clear that such sample restrictions are appropriate.

17. Duncan and Hill (1985).

18. The fraction of near-zero events for nontransfer labor income of the head of household was far higher, on the order of 3 to 4 percent of the sample. Thus including transfer income and spouse's income greatly reduces the incidence of near-zero events.

19. The information presented so far does not rule out the possibility that when a household's income drops to near zero, it stays near zero, or, less dramatically, it reverts to a much lower level than before. Investigation of this question, however, showed that when the sample excludes the self-employed and people experiencing changes in marital status, income typically recovers fully from near-zero events within three years, and mostly recovers within a year.

68

In light of these results, the model below will assume that the near-zero events are transitory shocks which occur with a probability of 0.5 percent each period.

Characterizing the Distributions When Income Does Not Go to Zero

The relative magnitudes of the transitory and permanent components of shocks when income does not go to zero now must be determined.²⁰ Recent work in the macroeconometric literature has illustrated the difficulties of decomposing univariate time series into transitory and permanent components; however, under sufficiently strong conditions it can be done. It turns out that the assumptions needed to solve the theoretical model below are also strong enough to generate a unique variance decomposition between permanent and transitory shocks. These assumptions are that the shocks to income are serially uncorrelated and independent at all leads and lags.

In addition to the sample limitations discussed above, the sample used for the transitory/permanent decomposition was further limited to heads of households whose marital status did not change; who did not experience a near-zero income event; and who were never self-employed. All three exclusions should reduce measured variability relative to true variability. Figure 2 presents a histogram of the distribution of the first differences in ln YL_{it} in this sample, along with a reasonably well-fitting normal distribution whose parameters were estimated from the data. Based on the results of figure 2, I will assume that both the transitory and the permanent shocks are lognormally distributed, producing the approximately lognormal distribution of annual innovations.²¹

Under the strong assumptions that have been made, the methodology for recovering the transitory and permanent components of the shocks is simple (and is essentially the same as the methodology used by Robert

20. See below for a discussion of measurement error as a possible source of some of the measured innovations in income.

21. To the extent that the normal distribution does not fit the histogram, it is principally due to the excessively fat tails in the true distribution. This means that assuming lognormality again errs on the side of minimizing the importance of uncertainty, because lognormality underrepresents draws in the lower tail of the income distribution, which are disproportionately influential in producing precautionary behavior.



Figure 2. Distribution of Annual Innovations in Log Family Noncapital Income Number of observations^a

Source: Author's calculations based on data from PSID as described in the text. a. Observations are equal to the number of households multiplied by the number of years in the sample period. b. Size of innovation is defined as the first differences in the log of detrended noncapital income.

Hall and Frederic Mishkin).²² Note that, defining $g = G - 1 \approx \ln G$, if income in future periods is compared to income today, it follows that

$$\ln P_{t+m} \approx \ln P_t + mg + \ln N_{t+1} + \ln N_{t+2} + \ldots + \ln N_{t+m}$$

$$\ln YL_{t+m} - \ln YL_t = (\ln P_{t+m} + \ln V_{t+m}) - (\ln P_t + \ln V_t)$$

$$\approx (\ln P_t + mg + \ln N_{t+1} + \ln N_{t+2} + \ldots + \ln V_{t+m})$$

$$- \ln P_t - \ln V_t$$

$$= mg + \ln N_{t+1} + \ldots + \ln N_{t+m} + \ln V_{t+m} - \ln V_t.$$

Because it is assumed that $\ln V_i \perp \ln V_j$ and $\ln N_i \perp \ln N_j$ for all $i \neq j$ and $\ln V_i \perp \ln N_j$ for all i, j, this implies that

(3)
$$\operatorname{var}(\ln YL_{t+m} - \ln YL_{t}) = m \,\sigma_{\ln N}^{2} + 2 \,\sigma_{\ln V}^{2}.$$

Because $\ln YL_{t+m} - \ln YL_t$ can be calculated for every household *i*, there is a sample on which to calculate var $(\ln YL_{t+m} - \ln YL_t)$ for any *t*

22. Hall and Mishkin (1982).

		Standard deviation of innovations ^b			
Definition of income and sample period ^a	Number of households	Transitory component	Permanent component		
Noncapital family income, 1976–85	834	0.163	0.126		
Total family income, 1976-85	835	0.164	0.128		
Total family income, 1968-85	398	0.155	0.113		
Labor income of head, 1968-85	395	0.171	0.105		

	Table 2.	Size Estimates	of the	Transitory	and Permanent	Components -	of Income	Shocks
--	----------	----------------	--------	------------	---------------	--------------	-----------	--------

Source: Author's calculations using PSID.

a. The sample is restricted to households whose "head" experienced no change in marital status, who was never self-employed, and who was between the ages of 25 and 63 over the entire period. All these restrictions reduce estimated income uncertainty. The sample is also restricted to households that did not experience a near-zero income event (that is, observations in which income was at least 10 percent of its average over the entire sample period). b. See text and appendix for methodology for calculating the standard deviation of transitory and permanent

shocks. Reported estimates are robust to MA(2) serial correlation for the transitory shocks.

and *m* that remain within the 1976–85 sample period. For any two different values of *m*, it is possible to solve two linear simultaneous equations for $\sigma_{\ln N}^2$ and $\sigma_{\ln V}^2$. The first section of the technical appendix shows how to generalize this technique to allow for MA(q) serial correlation in the transitory shocks.

Table 2 presents the estimates of $\sigma_{\ln N}^2$ and $\sigma_{\ln V}^2$ for a variety of samples and definitions of income. (The reported estimates are robust to MA(2) serial correlation in the transitory shocks; see the technical appendix.) For the preferred income definition, total noncapital income, the technique estimates the standard deviations of the annual innovations in both components to be approximately 0.15. Virtually identical results are obtained using total (capital and noncapital) income over the 1976– 85 period.

Interpreted literally, this would mean that every year, about onethird of consumers experience positive or negative shocks to their permanent income of greater than 15 percent. Similarly, each year about one-third (but not the same one-third) of consumers experience transitory shocks of more than 15 percent in each year. These estimates of the uncertainty of income may seem surprisingly high, but they are consistent with estimates obtained by many other researchers.²³ Some authors have speculated that measurement error may account for most of the annual innovations in income, but several empirical studies have con-

23. The raw estimates of the annual innovations in labor income of the household head, for instance, are quite close to estimates obtained by Abowd and Card (1989).

cluded that most of the variation in measured earnings corresponds to true variation. John Bound and Alan B. Krueger, for instance, conclude that around 70 percent of the variation in their first-differenced sample corresponds to true variation and not to measurement error.²⁴ Another strike against the view that most of the variation in income is measurement error is the large size of the estimated permanent component of the annual innovation. It is hard to see why measurement error would have a unit root.

In light of these results, the income distribution process assumed for the rest of the paper will be as follows. In each period, there is a 0.5 percent chance of a transitory shock that causes income to go to zero in that period. If income is not zero, shocks to both transitory and permanent income are lognormally distributed with standard deviations $\sigma_{\ln V} =$ $\sigma_{\ln N} = 0.10$. These numbers are substantially lower than suggested by the data in order to compensate for whatever measurement error may exist, and also because it seems better to understate, rather than overstate, income uncertainty. If, as below, precautionary saving effects are still found to be important, the true effects are probably larger than those estimated using conservative measures of uncertainty.

The Basic Model

This section of the paper summarizes the characteristics of the bufferstock model originally developed in my other 1992 paper.²⁵ After formally specifying the model and choosing parameter values, I show why buffer-stock behavior emerges from this framework. I then consider how changing various parameters would affect the long-run target for wealth. Next I provide an example of the short-term dynamic reaction of consumption, saving, and wealth to a change in the long-run target. Finally, I discuss how the model would differ if explicit liquidity constraints were imposed.

Parameterizing and Solving the Model

Consider a standard intertemporal consumption model:

(4)
$$\max E \sum_{t=0}^{T} \beta^{t} u(C_{t}),$$

24. Bound and Krueger (1989).

^{25.} Carroll (1992).

such that

$$W_{t+1} = R(W_t + YL_t - C_t)$$
$$YL_t = P_t V_t$$
$$P_{t+1} = G P_t N_{t+1},$$

where YL, P, V, and N are, as previously defined, current and permanent labor income and their shocks; G = 1 + g, where G is the growth factor and g is the growth rate; W is net wealth; R = 1 + r, where r is the interest rate and R is the interest factor; $\beta = 1/(1 + \delta)$ is the discount factor, where δ is the discount rate; and C is consumption. The standard Constant Relative Risk Aversion (CRRA) utility function is of the form $u(C) = C^{1-\rho}/(1 - \rho)$ where ρ is the coefficient of relative risk aversion. I do not explore here the implications of a model like that of Larry Epstein and Stanley Zin in which the intertemporal elasticity of substitution is allowed to differ from the inverse of the coefficient of relative risk aversion.²⁶

It turns out that the optimal consumption rule in a model with both transitory and permanent shocks can be written as a relationship between the ratio of consumption to permanent income, C/P, and the ratio of gross wealth X to permanent income, X/P, where gross wealth X is defined as assets plus current income: $X_t = W_t + YL_t$. If lower case variables denote the upper case counterpart normalized by permanent income, this means that the optimal consumption rule can be written as $c(x_t)$. That is, the optimal consumption ratio is a function of the gross wealth ratio and nothing else.²⁷

The model can now be solved for the optimal consumption rule for arbitrary sets of parameter values. The uncertainty parameters estimated in the previous section are p (zero income) = 0.5 percent, and $\sigma_{\ln V} = \sigma_{\ln N} = 0.10$. This leaves the discount factor β , the interest factor R, the income growth factor G, and the coefficient of relative risk aversion ρ still to be chosen.

Although many empirical estimates have placed the value of the coefficient of relative risk aversion ρ at 5 or greater, such a value produces unreasonably large amounts of precautionary saving, even for consumers who have a 10 percent annual discount rate.²⁸ After some experi-

- 27. For a proof, see the second part of the appendix or my 1992 paper (Carroll, 1992).
- 28. For estimates from microeconomic data, see Mankiw and Zeldes (1990).

^{26.} Epstein and Zin (1989).

mentation, I settled on a value of 3 for the simulations reported in this paper.

In order to generate buffer-stock saving behavior, it is necessary that consumers be impatient in the sense that if they faced no income uncertainty, they would want to borrow against their future income in order to consume more today. Making the usual approximations that r = $\ln R$, $g \approx \ln G$, and $-\delta \approx \ln \beta$, it turns out that the condition that is required for consumers to exhibit impatience in this sense is ρ^{-1} $(r - \delta) < g^{29}$ To see why this is a necessary condition, note first that if income were perfectly certain, the growth rate of consumption would be approximately $\rho^{-1} (r - \delta)$.³⁰ Now consider a consumer with zero wealth and recall that the present discounted value (PDV) of consumption equals the PDV of income. If consumption is growing more slowly than income but has the same PDV, the level of consumption must be higher than the level of income. (If consumption were less than income and growing more slowly, it clearly would have a lower (PDV.) Thus if income were certain and $\rho^{-1}(r-\delta) < g$, consumers would wish to spend more than income; that is, they would wish to borrow.

Impatience is necessary because it limits the size of the net wealth that consumers will be willing to accumulate. The qualitative results from the buffer-stock model are similar whether consumers wish to borrow because they have high discount rates or because they have high expected growth rates of income, but the base assumption will be of a rather high discount rate of 10 percent and a relatively low expected income growth rate of 2 percent annually. Limited results will also be presented for other values of these parameters.

The solution method, described in the appendix, involves solving for the optimal consumption rule in the last period of life, then the optimal rule in the second-to-last period, then the third-to-last period, and so

29. $\rho^{-1}(r-\delta) < g$ is a necessary condition but, I believe, not a sufficient condition for generating buffer-stock behavior. I believe that the necessary and sufficient condition is that $\rho^{-1}(r-\delta) + \rho/2 \sigma^2_{\ln N} < g$; this is the condition that Deaton (1991) derives for his model, which is similar to the model presented here, but which contains explicit liquidity constraints. However, I have been unable to prove rigorously that this condition is sufficient to guarantee the emergence of buffer-stock saving behavior.

30. To derive this formally, note that the Euler equation is $u'(C_t) = \beta R u'(C_{t+1})$, implying $(C_t)^{-\rho} = \beta R (C_{t+1})^{-\rho}$, implying $C_{t+1}/C_t = (\beta R)^{1/\rho}$. Taking logs of both sides gives $\ln C_{t+1} - \ln C_t = \rho^{-1} (\ln R + \ln \beta) \approx \rho^{-1} (r - \delta)$.



Figure 3. Optimal Consumption Rules According to Years of Life Remaining

forth, until the difference between successive consumption rules is small enough that the consumption rule can be said to have converged. Optimal consumption rules at selected ages under the base parameter assumptions are presented in figure 3. Note that, under the specific set of assumptions and parameter values chosen here, the optimal consumption rules converge rapidly: the difference between the consumption rules ten years before the end of life and forty-nine years before the end of life is almost invisible. What this means is that the behavior of consumers more than ten years from the end of life is virtually the same whether they have eleven years of life remaining or fifty.³¹

The fact that the consumption rules converge so quickly appears to support Hall's assertion that there is little practical difference between the implications of a "life-cycle" model in which consumers face death in a finite number of periods and the implications of a "permanent-

Source: Author's calculations as described in the text. The end of life is T. The ratio of consumption to permanent income is c. The gross wealth ratio, x, is the ratio of gross wealth to permanent income.

^{31.} See Carroll (1992) for more results from the model over a finite lifetime.

income" model in which there is an infinite horizon.³² To some extent, however, this conclusion depends on the parameter values chosen here. A particularly interesting point is that if consumers face an income profile with rapid growth early in life followed by slower growth (or declines) later in life, it is possible that they might switch from buffer-stock saving behavior when young to more traditional life-cycle saving behavior as retirement approaches. Intuitively, this is because when they are young they face a high income growth rate so they are impatient in the $\rho^{-1} (r - \delta) < g$ sense, but as g falls with age, eventually they reach a point where they are no longer impatient in this sense. In my 1992 paper, I explore this possibility in greater detail, but for the purposes of this paper I will only consider "permanent-income" consumers who face the same consumption rule and income process every year into the indefinite future.³³

Why Does Buffer-Stock Saving Behavior Arise in this Model?

The intuition for how the model generates buffer-stock saving behavior is given by considering the Euler equation for consumption growth. If shocks to consumption are lognormally distributed, it can be shown that consumption will grow according to

(5)
$$\Delta \ln C_{t+1} \approx \rho^{-1} (r - \delta) + \frac{1}{2} \rho E_t \operatorname{var} (\Delta \ln C_{t+1}) + e_{t+1}$$

The key insight here is that the expected variance of consumption growth, $E_t \operatorname{var}(\Delta \ln C_{t+1})$, will be negatively related to wealth. Intuitively, consumers with less wealth have less ability to buffer consumption against income shocks; thus they have higher $E_t \operatorname{var}(\Delta \ln C_{t+1})$ and faster consumption growth. (The third part of the technical appendix derives this result formally.) The growth rate of consumption is high when wealth is small because the *level* of consumption is being depressed by precautionary saving. Kimball showed that for utility functions such as the one considered here, precautionary saving declines as wealth increases.³⁴ Therefore, over time, as precautionary saving adds to wealth,

32. Hall (1978).

33. Carroll (1992).

34. Kimball (1990a, b). This is true for utility functions that exhibit decreasing absolute risk aversion.

consumption will become less depressed. It is this lessening of precautionary saving as wealth increases that generates the extra growth in consumption when $E_t \operatorname{var} (\Delta \ln C_{t+1})$ is high.

Note that in this model lagged income shocks will affect current consumption growth because lagged income shocks surely affect current wealth. Thus the legion of tests following Hall,³⁵ using both macroeconomic³⁶ and microeconomic³⁷ data that claim to reject the LC/PIH because lagged income growth predicts current consumption growth, do *not* constitute rejection of general versions of the LC/PIH framework that incorporate precautionary saving motives. In econometric terms, there is an omitted variable, the conditional variance of consumption growth, which ought to be correlated with the lagged change in income.³⁸

The obvious argument is that a positive shock to income in the last period implies higher wealth both in the last period and in this period, which in turn implies a higher level of consumption in both periods and a lower growth rate of consumption between periods. This would predict that regressions of consumption growth on lagged income changes would find a negative coefficient, as Zeldes found.³⁹ However, regressions using macroeconomic data typically find a positive coefficient on lagged income changes. This could be justified if rapid income growth occurs in periods when wealth is low. For instance, the periods of fastest income growth might be in the early stages of recovery when buffer stocks have been depleted by recession. In this case, the model might predict a positive coefficient on lagged income changes.

Figure 4 attempts to provide some intuition about how the bufferstock model works by showing the relationship between the gross

- 35. Hall (1978).
- 36. See, for example, Flavin (1981).
- 37. See, for example, Zeldes (1989b).

38. Note that the relevant variable is the conditional variance of consumption growth at the *household* level. There is little reason to expect that this will be related to the conditional variance of aggregate consumption growth. Hence, there is little reason to expect success from simple attempts to remedy previous tests of the Euler equation by, say, estimating an ARCH process for the variance of aggregate consumption growth and including predicted aggregate variance on the right hand side. Some more direct measure of household consumption uncertainty is necessary. In the fourth section of this paper, I argue that unemployment expectations may be an appropriate proxy.

39. Zeldes (1989b).

Figure 4. The Buffer-Stock Relationship between Growth Rates and the Gross Wealth Ratio

Growth rates of income and consumption



Source: Author's calculations as described in the text. Arrows represent convergence toward the "target" wealth ratio. Gross wealth ratio is the ratio of gross wealth to permanent income. The minimum possible amount of gross wealth is h. The expected consumption growth next period is $E_t \Delta \ln C_{t+1}$. The target wealth on hand is x^* . The perfect certainty growth rate of consumption is $\rho^{-1}(r - \delta)$. The growth rate of income is g.

wealth ratio, expected consumption growth, and income growth.⁴⁰ Consumers are impatient in the sense described above: $g > \rho^{-1} (r - \delta)$. The line with arrows traces out the relation between expected consumption growth in the next period, $E_t \Delta \ln C_{t+1}$, and gross wealth in the current period, x_t . As x_t approaches its minimum possible level <u>h</u>, expected consumption growth approaches infinity. This is because as gross wealth x_t approaches <u>h</u>, C_t must approach zero (see the next paragraph for arguments demonstrating this point). Because the expected level of C_{t+1} is positive even if the consumer enters period t + 1 with no assets, it follows that as x_t approaches <u>h</u>, $E_t \Delta \ln C_{t+1}$ approaches infinity. On the other hand, as x_t approaches infinity, uncertainty about future labor

40. Figure 4, and much of the following discussion, draw upon work in my 1992 paper (Carroll, 1992).

income becomes essentially irrelevant to consumption because expected future labor income becomes infinitesimal in relation to current wealth. Thus as x_i approaches infinity, the consumption growth rate approaches the growth rate under certainty, $\rho^{-1}(r - \delta)$. Because everything in the model is continuous and monotonic, the expected consumption growth rate will cross the income growth rate curve at one point. The gross wealth ratio at this point will be called x^* , the "target" gross wealth ratio.

The vertical axis in figure 4 is drawn at a place labeled h. This is meant to signify (the negative of) the minimum possible present discounted value of future labor income. For example, if in each period the minimum possible value of labor income were some amount Y, then the worst possible outcome would be for income to equal Y in every future period. If that were to occur, the PDV of labor income would be given by Y/r, so h would be -Y/r. Under the parameterization described earlier in the paper, a chance exists that labor income could go to zero in each future year, so in this case, h = Y = 0. This implies that consumers will never borrow on net, even though there are no formal liquidity constraints.⁴¹ To see why, consider a consumer in the second-to-last period.⁴² If that consumer saved nothing and entered the last period with zero wealth, there would be some possibility of earning zero income in the last period, in which case consumption would be zero. Because consumption of zero yields negative infinite utility, no consumer will ever allow himself to enter the last period of life with zero wealth-consumers will always save at least a little bit to ensure positive consumption. Now consider a consumer two periods from the end of life. A chance of earning zero income exists in both the last and the second-tolast periods. The consumer in the third-to-last period will therefore never allow himself to enter the second-to-last period with zero assets. This logic can be applied recursively back to the first period of life, so consumers will never borrow.

Now suppose income cannot go all the way to zero. Perhaps social insurance, private charities, or other sources would provide income of at least Y > 0, even in the event of, say, total disability. In the absence

42. The following repeats arguments in Zeldes (1989a) and Schechtman (1976).

^{41.} To be precise, net worth will never be negative. This therefore does not rule out mortgages, auto loans, and other loans used to purchase assets whose value exceeds the debt used to purchase them.

of liquidity constraints, it is easy to show that the implications of such a model would be essentially identical to the model just described, except that the minimum possible level of gross wealth \underline{h} (the vertical axis in figure 4) would no longer be zero, but would instead equal $-\underline{Y}/r$. In this case, wealth could be negative, but never less than $-\underline{Y}/r$. Thus the existence of a positive lower bound on income does not change the qualitative characteristics of the model, unless an additional assumption of liquidity constraints is made. (See the final part of this section for a discussion of how liquidity constraints change the model.)

Returning to figure 4, the gross wealth ratio x^* is a target in the sense that, if actual gross wealth is below x^* , the consumer will spend an amount small enough so that gross wealth will be expected to increase; however, if actual gross wealth is greater than x^* , the consumer will spend enough so that expected gross wealth next period will decline (hence the arrows on the expected growth rate curve). Section D of the appendix proves this for a simplified case. Section E of the appendix proves that, if gross wealth is at the target ratio, the expected growth rate of consumption is approximately equal to the growth rate of income.⁴³

It is easy to use figure 4 to see how the model's results differ under alternative parameter values. For instance, increasing g, the growth rate of future income, will decrease the target wealth stock by shifting the intersection with the $E_t \Delta \ln C_{t+1}$ curve left. This is a standard effect; higher future income results in higher current consumption, hence lower saving and lower wealth. Increasing the discount rate will shift ρ^{-1} $(r - \delta)$ down; therefore the $E_t \Delta \ln C_{t+1}$ curve will also shift down, decreasing target wealth. Increasing the interest rate will shift $\rho^{-1} (r - \delta)$ up, nudging $E_t \Delta \ln C_{t+1}$ up and increasing target wealth. Increasing the degree of uncertainty in income will increase the variance of consumption growth for any level of wealth, directly shifting the $E_t \Delta \ln C_{t+1}$ curve up and increasing target wealth. Increasing the coefficient of relative risk aversion ρ will have two effects. It will shift the $E_t \Delta \ln C_{t+1}$ curve up

^{43.} Actual gross wealth for an individual consumer, of course, will be subject to random shocks, and thus for any individual, x_{t+1} may end up being either closer to x^* than x_t was, or farther away, depending on next period's random draws from the uncertainty distributions. Convergence toward x^* happens in the sense that, if a large number of consumers with the same value of x_t were considered, *average* x_{t+1} next period would be closer to x^* than x_t was.

as $[\rho \operatorname{var} (\Delta \ln C_{r+1})]$ increases, tending to increase wealth (this is the direct effect of increased risk aversion). It also will shift the $\rho^{-1}(r-\delta)$ curve down, thus tending to reduce wealth (this is the effect of a lower intertemporal elasticity of substitution).

Figure 4 and the proofs in the appendix relate consumption to the gross wealth ratio x^* because the consumption ratio is a direct function of the gross wealth ratio. Most of the remaining analysis of the paper, however, is more easily performed using the net wealth ratio w = W/P. In particular, it will be useful to define the target net wealth-to-income ratio w^* as $w^* = x^* - 1$. This is justified by the fact that gross wealth is X = W + YL; dividing through by P, x = w + V. Because the expected value of the transitory factor V is one (by definition), it is sensible to define the target net wealth ratio as $w^* = x^* - 1$.

One reason it is useful to define w^* is that a simple relationship exists between w^* and the personal saving rate. If $w^* = W^*/P$ is to remain constant, net wealth W must be growing at the same rate g as permanent income P. If the interest rate were zero, for W^* to grow at rate g, the consumer would have to save an amount gW^* . Dividing the amount of saving by the level of permanent labor income, P, yields the personal saving rate out of labor income at the target net wealth ratio, so that $s^* \approx gw^*$. Because the net wealth ratio will tend toward w^* in a bufferstock model, the personal saving rate will tend to approach gw^* over time. The last section of the appendix shows that even with a positive interest rate, the personal saving rate at the target net wealth-to-income ratio will be well approximated by $s^* \approx gw^*$.

Table 3 gives some quantitative results on how w^* differs under a variety of parameter values. Starting with the sensitivity of w^* to the probability of zero-income events, the first column of table 3 presents values for w^* as p (zero income) varies from 0.1 percent to 1 percent annually.⁴⁴ At p = 0.1 percent annually, the target buffer stock is 26 percent of annual income; at p = 0.5 percent, it is 44 percent of income; and at p = 1 percent, it is 56 percent of income, as seen in the last column of table 3. Very bad events clearly have a powerful effect on saving behavior, even

44. If the distributions of the transitory and permanent shocks remained the same as the probability of zero-income events changed, mean expected income would change. This effect was offset by increasing the mean of the transitory distribution by an amount such that expected income in the next period is the same for all values of the probability of zero-income events. when the events are as rare as those considered here. This strongly suggests that it is important to include some treatment of unemployment in any buffer-stock type model of precautionary saving.⁴⁵

The next two columns of the table show how w^* differs under values of $\sigma_{\ln V}$ and $\sigma_{\ln N}$ ranging from 0 percent annually to 15 percent annually.⁴⁶ In comparison with the effect of changing *p* (zero income), changing the variance of either transitory or permanent income shocks has comparatively modest effects.⁴⁷

The fourth column of the table presents the results as the real interest rate is increased from 0 percent to 2 percent to 4 percent annually. The target net wealth-to-income ratio is barely affected by the interest rate, moving from 0.44 at r = 0 percent to 0.53 at r = 4 percent annually. Note that this implies that the long-term elasticity of the personal saving rate with respect to the interest rate is approximately zero. In contrast, Summers showed that the life-cycle model implies a strongly positive interest elasticity of saving;⁴⁸ the permanent income model implies an even stronger positive elasticity. This prediction has been something of an embarrassment for the LC/PIH model because empirical estimates of the interest elasticity of saving have typically been quite small.⁴⁹ The empirical failure has been particularly acute over the last decade, when real interest rates have been very high and personal saving has been very low.

Results are also presented for different values of the discount rate, the growth rate of income, and the coefficient of relative risk aversion ρ . A choice of $\rho = 1$ (log utility) generates a target net wealth-to-income

45. The argument could be made that the treatment here is excessively stylized. A more realistic model might be a compromise between Deaton's model (1991) and the model described here. If income could drop to, say, 40 or 50 percent of its previous level, but did so with a probability of, say, 5 percent (to approximate the true incidence of job loss), one might well obtain results similar to those produced here. However, I have not performed such simulations.

46. The distributions for the transitory and permanent shocks were truncated at three standard deviations from the mean.

47. Because these transitory and permanent shocks are lognormal rather than normal distributions, increasing the variance slightly increases the mean level of the variable. This increase in the mean level was sterilized by reducing the mean of the lognormal distributions by an amount sufficient to keep the mean levels the same as before the increase in variance.

49. Carroll and Summers (1991).

82

^{48.} Summers (1981).

Probability of zero income, p (0.5) ^b	$\begin{array}{c} Standard\\ deviation\\ of\\ transitory\\ shocks,\\ \sigma_{\ln V}\\ (0.10) \end{array}$	Standard deviation of permanent shocks, o _{InN} (0.10)	Real interest rate, r (0.0) ^c	Growth rate, g (2.0) ^d	Discount rate, δ (10.0)°	Coefficient of relative risk aversion, p (3.0)	Target wealth-to- income ratio
0.1							0.26
0.5							0.44
1.0							0.56
	0.00						0.42
	0.10						0.44
	0.15						0.46
		0.00					0.39
		0.10					0.44
		0.15					0.57
			0.0				0.44
			2.0				0.48
			4.0				0.53
				0.0			0.60
				2.0			0.44
				4.0			0.37
					5.0		0.54
					10.0		0.44
					15.0		0.38
						1	0.06
						3	0.44
						5	0.88

Table 3. Target Net Wealth Ratios with Alternative Parameter Values Parameter value^a

Source: Author's calculations as described in the text.

a. Each column keeps all parameters fixed at their base values, except the parameter in the column heading, which is allowed to vary. When all parameters are at their base values, the target net wealth-to-income ratio is 0.44. The base value for each parameter is shown in parentheses below the column heading.

b. Probabilities of zero income are shown as a percent.

c. Real interest rates are shown as a percent.

d. Growth rates are shown as a percent.

e. Discount rates are shown as a percent.

ratio of only 6 percent of annual income, while $\rho = 5$ increases the target to 88 percent of income. Of course, it is not surprising to find that the amount of precautionary saving is very sensitive to the degree of consumers' risk aversion. Varying the other parameters generally moves the target ratio from 35 percent to 60 percent of annual income.

Informal evidence that this is a reasonable characterization of behavior comes from personal financial planning guides. These guides commonly have passages that advise consumers to maintain a buffer stock of assets. The following excerpt from a 1989 Touche Ross financial planning guide provides an example:

It is generally held that your liquid assets should roughly equal four to six months' employment income. If you are in an unstable employment situation . . . the amount should probably be greater. [Your] need for liquidity is determined by the predictability of your cash income and expenditures, by your employment security, and by your investment strategy.⁵⁰

The Dynamic Reaction to a Change in Target Wealth: An Example

The figures in table 3 showed how the target net wealth ratio depends on the model's parameter values, and figure 4 and the discussion showed that in the long run, gross wealth will converge to the target. What has not been discussed yet, however, is the short-run dynamics of consumption, saving, and wealth when parameters change. For example, suppose the probability of unemployment and zero-income events increased from 0.5 percent to 1 percent per year.⁵¹ Table 3 showed that the target net wealth/income ratio will increase from 0.44 to 0.56. But what happens to consumption, saving, and wealth immediately after the change? The first consequence, illustrated in figure 5, is that the new optimal consumption rule is lower than the old one, indicating that, at any given level of gross wealth, in the new more uncertain environment, the consumer will save more than before.

To trace out the dynamic behavior of consumption, saving, and wealth, I assume that the consumer begins holding the optimal gross wealth ratio x_1^* implied by the old unemployment expectations p = 0.5percent. I then assume that the consumer just happens to experience the mean values, N = V = 1, of the uncertainty variables for the next few years of life (although this does not affect his expectations about the fu-

50. Touche Ross (1989, p. 10).

51. This exercise is not strictly valid because the model is solved under the assumption that the probability of zero-income events is fixed, yet in the exercise here I assume that the probability changes from 0.5 percent to 1 percent. The fully correct way to do this exercise would be to specify a Markov transition matrix for the probability of a zero-income event and then to trace out the reaction when a transition occurs. The results would presumably be qualitatively similar, but the magnitude of the effects described above would depend on the properties of the Markov transition matrix.



Consumption ratio, c



Source: Author's calculations as described in the text. The probability of income falling to zero is given by p.

ture distributions of these variables).⁵² Applying the consumption rule recursively traces out the dynamic paths of consumption, wealth, and saving.

The results can be seen in figure 6. The change in expectations is assumed to occur between years 5 and 6. Leading up to year 5, the consumer was in equilibrium with constant consumption, net wealth, and saving ratios. In year 6, the first year of the new regime, consumption drops sharply and the saving rate rises. Net wealth therefore begins to increase; as it does so, consumption begins to recover toward its longrun target level. After its initial surge, the saving rate declines steadily toward its new, higher target level. In the end, the consumption ratio is slightly less than before the shock, while the saving rate and the net

52. The completely correct way to do this experiment would be to derive the steadystate distribution for x under the first set of assumptions, and then calculate the evolution of this distribution over time subject to the new consumption rule and the given distributions for uncertainty. Simulations of this kind in my 1992 paper (Carroll, 1992) show that the procedure described here provides a good approximation to the behavior of the means of the variables in the full simulation model.



Figure 6. The Consumption Ratio, the Personal Saving Rate, and the Wealth Ratio The ratio of consumption to personal income, c

Source: Author's calculations as described in the text.

wealth ratio are both higher. Across steady states, the change in the personal saving rate will be given by comparing gw_1^* to gw_2^* , where w_1^* indicates the target net wealth ratio before the change and w_2^* is the ratio afterwards. So long as w_1^* is positive, if the growth rate of income is positive, the new target personal saving rate will be higher. On the other hand, the target saving rate gw_2^* will be lower than the rate during the transition period, as wealth is being built up from w_1^* to w_2^* .

How Would the Model Differ with Explicit Liquidity Constraints?

Despite the evidence of the second section, the assumption that income can go to zero and potentially stay there for the remainder of life seems somewhat artificial. Arguments above demonstrated, however, that simply assuming that income has a lower bound above zero does not change the model qualitatively; it just makes h—the vertical axis in figure 4-negative rather than zero. For lower bounds on income to have any effect other than a shift of axes, explicit borrowing constraints must also be present and relevant. (To be relevant, constraints must limit wealth to be somewhere above h, because consumers would not wish to borrow more than h anyway.) It would be useful to treat such constraints explicitly because it would then be possible to explore the quantitative consequences of changing the constraints. However, the qualitative consequences are intuitive and can be summarized as follows. Imagine a model like Deaton's in which consumers are impatient, have a positive minimum future labor income (h < 0), but cannot borrow at all (w = 0). Suppose, further, that on average they hold a buffer stock of size w_1^* . If they are suddenly allowed to borrow, say, up to amount w < 0, they will eventually reach a new steady-state equilibrium in which they on average hold a target wealth ratio, w_2^* , that is less than w_1^* . Thus in a buffer-stock model with relevant borrowing constraints, loosening the constraints should result in behavior qualitatively identical to the effect of increasing the minimum guaranteed future income in a model without liquidity constraints. As long as average wealth is positive both before and after the loosening of constraints, the dynamic behavior of saving would also be the same: a sharp drop at first, then the partial recovery to stabilize at a new rate gw_2^* that is lower than the initial rate gw_1^* .

It would be interesting to analyze liquidity constraints quantitatively as well as qualitatively, but that exercise is beyond the scope of this paper. Some results can be found in Deaton's 1991 paper. In particular, Deaton finds that, under his parameter values, consumers often hold exactly zero assets.⁵³ The rest of the time they hold a small buffer stock, as in the model presented here. However, compared to the assumptions here, Deaton's income process is quite benign.⁵⁴ Most importantly, there is nothing in his model that resembles a spell of unemployment. But the evidence of the second section of this paper, as well as casual experience, indicate that in the real world, outcomes occur that are much worse than any outcome Deaton considers.

Of course, it is possible that these bad events are so rare that Deaton did not miss much by ignoring them. The only way to know is to solve the model numerically under a range of values for the probability of unemployment to see what the consequences are. But table 3 showed that the size of the target buffer stock, w^* , was quite sensitive to the probability of zero-income events, strongly suggesting that the results from a model that ignores the possibility of unemployment are likely to be substantially different from the results when unemployment is explicitly treated.

Dynan's Empirical Test

Before I present my own empirical evidence, a brief discussion of an interesting recent paper by Karen Dynan is in order.⁵⁵ Using data from the 1985 Consumer Expenditure Survey, Dynan calculates quarterly consumption growth and the variance of quarterly consumption growth for a sample of households, and estimates an equation like equation 5, using instrumental variables for the E_t var ($\Delta \ln C_{t+1}$) term, where the instruments are demographic variables such as age, occupation, and education. She obtains very low estimates of ρ , and concludes that the certainty equivalence model cannot be rejected using her data.

I can think of two main reasons why a test like Dynan's might obtain low or zero estimates of ρ , even if most consumers have a value of ρ in

54. In the published version of his paper, I could not find an indication of the bounds of his income distribution process; but if memory serves, an earlier draft stated that the bounds were two standard deviations of income. Because the standard deviation of income is 0.1, this gives a lower bound of a 20 percent drop in income.

55. Dynan (1991).

^{53.} Deaton (1991).

the range assumed in this paper. The first is a subtle theoretical point; imagine estimating Dynan's model on a population of buffer-stock consumers with exactly the same expected income growth rate g, but with some variation across demographic groups (no matter how slight) in ρ , r, or δ . Because the consumption growth rate converges to the income growth rate for buffer-stock consumers, the expectation of consumption growth for each consumer will be g, yet consumers in different groups will have different values of $\rho^{-1}(r - \delta)$, so that consumers in each group will be converging toward a different target buffer-stock and a different target value for E_t var ($\Delta \ln C_{t+1}$). Assuming Dynan's instruments would correctly identify the target value for E_t var ($\Delta \ln C_{t+1}$) for each demographic group, her instrumental variables test should estimate a value of ρ exactly equal to zero, because expected consumption growth would be the same, g, for every consumer; however, E_t var ($\Delta \ln C_{t+1}$) would vary by demographic group.

The second objection to Dynan's test is empirical. Actually testing an Euler equation like 5 on quarterly data, at the household level, requires heroic assumptions about the quality and meaning of the data. In practice, the variability of consumption growth at frequencies as short as one quarter is likely to be mostly unrelated to uncertainty. The mere fact that the quarterly standard deviation of consumption growth is about 20 percent in her sample should give pause; this number seems far too high to represent quarterly uncertainty. It seems much more likely that almost all of the quarterly variation has nothing to do with uncertainty. For instance, the variance of quarterly consumption expenditures is likely to depend on whether vacations are long or short, and the length of vacations is likely to be correlated with occupation. Alternatively, people with more education might be more likely to send their children to expensive schools, causing large fluctuations in quarterly spending between the summer and the school year. Instruments would discover higher var ($\Delta \ln C$) for occupations with long vacations, or for more educated people, but this would be unrelated to uncertainty and would therefore not correspond to higher consumption growth. Thus the estimate of ρ would be biased toward zero.

Dynan's paper is a valuable contribution, but I suspect that her tests simply demand too much of her data. The buffer-stock model delivers a wide variety of testable predictions, and my own instinct is that tests using lower frequency data, and tests based on the level of saving rather than the growth rate of consumption, are likely to be more robust than Euler equation tests such as Dynan's using quarterly household-level consumption. I turn now to some macroeconomic evidence along these lines.

The Buffer-Stock Model and Some Empirical Evidence

This section argues that the buffer-stock model provides a better explanation for a variety of empirical findings than simple versions of the LC/PIH model, a simple Keynesian model, or a hybrid like that proposed by Campbell and Mankiw.⁵⁶ At the outset, it is important to be precise about what is meant by each model. What this section refers to as the simple LC/PIH model is any model in which the principal purpose of saving is to smooth consumption over large, low-frequency, predictable fluctuations in income. The simplest such models are the various certainty-equivalence models (CEQ), although models with constant absolute risk aversion would also generally fall into this category. The simple Keynesian model will be taken to be a model of the form C = $\alpha_0 + \alpha_1 Y$, where consumption has a positive intercept α_0 and there is a constant marginal propensity to consume α_1 that is close to one (perhaps 0.9). The Campbell-Mankiw model blends these two models by assuming that half of income goes to LC/PIH consumers and half goes to Keynesian consumers with parameters $\alpha_0 = 0$ and $\alpha_1 = 1$.

This section will consider whether these models are consistent with three categories of macroeconomic evidence. First is some evidence that Lawrence Summers and I developed.⁵⁷ We found that aggregate consumption growth and income growth appear to be closely linked over periods of a few years or longer. In contrast to the LC/PIH model, the buffer-stock model provides a simple explanation: if there is a target net-to-wealth income ratio, consumption growth and income growth *must* converge; otherwise wealth would explode up or down. The second category is evidence that the personal saving rate depends strongly on the unemployment rate and expectations about future unemployment, even controlling for expectations about future income growth.

- 56. Campbell and Mankiw (1989).
- 57. Carroll and Summers (1991).

Such results strongly suggest the presence of a precautionary saving motive. The final category of evidence comes from estimating Euler equations for consumption growth. I find that the growth rate of consumption is slower in periods when unemployment fears are high. Again, the certainty equivalence model contains no explanation for this. However, the buffer-stock model also has some difficulty explaining the results. The combined results of this section suggest that when consumers become more pessimistic about the prospects of unemployment, they adjust their consumption downward. However, the adjustment is gradual rather than instantaneous.

The Consumption/Income Parallel

Figure 7 illustrates one of the most robust findings of my 1991 paper with Lawrence Summers: across countries, and within the same country over time, over periods of a few years or longer, the growth rate of consumption tends to be very close to the growth rate of income.⁵⁸ We find near-equality of consumption growth and income growth rates over periods as short as five years.

If aggregate consumption is modeled as resulting from a representative agent's behavior, the Euler equation 5, reproduced here for convenience, will apply:

(5)
$$\Delta \ln C_{t+1} = \rho^{-1} (r - \delta) + \frac{1}{2} \rho E_t \operatorname{var} (\Delta \ln C_{t+1}) + e_{t+1}.$$

The growth rate of income does not enter directly anywhere in equation 5. If we assume that the expected variance of consumption growth E_t var ($\Delta \ln C_{t+1}$) is constant, equation 5 provides no direct explanation for the consumption/income parallel of figure 7.

In the long run, of course, countries cannot violate their budget constraints, and under suitable conditions it can be shown that consumption growth will eventually converge to income growth even in an economy where a facsimile of equation 5 holds and E_t var ($\Delta \ln C_{t+1}$) is constant (perhaps zero). The mechanism for this is adjustment of the capital stock. If consumption is too high and growing more slowly than income, the capital stock will be declining. As the capital stock declines, the in-

58. Carroll and Summers (1991).



Figure 7. Income and Consumption Growth in OECD Countries, Various Periods, 1960–85

terest rate increases, and, through the first term in equation 5, as the interest rate increases, the growth rate of consumption increases toward the growth rate of income. Conversely, if consumption is low and growing faster than income, the capital stock will be increasing, driving interest rates down and reducing the consumption growth rate. Either way, the steady state is eventually achieved with consumption growth equal to income growth.

Summers and I argue, however, that this mechanism cannot explain why consumption and income growth rates are close at frequencies of

92

five years or less.⁵⁹ Adjusting the capital stock should take decades, not years. A similar objection applies to life-cycle explanations of the consumption/income parallel that rely on intergenerational productivity differences to produce an equilibrium where consumption growth equals income growth. Such models take at least a full generation to equate consumption growth with income growth.

Of course, if virtually all of consumption is done by Keynesian consumers with α_1 near one, the consumption/income parallel is no mystery. But the Keynesian model performs poorly at higher frequencies, where consumption and income growth are much less closely related, and is completely unable to explain microeconomic data on consumption and income.

The proposed solution to this puzzle comes from recognizing, as discussed in the third section above, that the variance of consumption growth depends importantly on the level of wealth. If most consumption is done by impatient buffer-stock consumers, then, as shown in figure 4, consumption growth will converge to income growth. Intuitively, if consumption is too low and growing too fast, the consumer must be saving. This saving boosts the wealth stock, which reduces the variance of consumption growth, which reduces the expected growth rate of consumption down toward the growth rate of income. If consumption is too high and growing too slowly, the consumer will be dissaving; as wealth falls, the expected variance of consumption growth rises, increasing the growth rate of consumption up toward the growth rate of income.

The equality of income and consumption growth rates in the bufferstock framework holds at the target net wealth stock w^* , but this prompts the question of whether consumption growth will approach income growth *rapidly* enough to explain the consumption/income parallel. Table 4 addresses this question by examining the transition path after the growth rate of permanent income changes in the model. The exact experiment is as follows. The optimal consumption rule is derived under the base assumptions described above, first with an expected growth rate g of 2 percent, and then with a growth rate of 0 percent annually. Consider a consumer in year T - 1 who was holding exactly his target buffer stock of assets, so that his expected growth rate of consumption in year T was 2 percent. Now suppose that this consumer

59. Carroll and Summers (1991).

Year	Income growth rate	Consumption growth rate	
T-1	2	2.00	
Т	0	-2.85	
T+1	0	0.99	
T+2	0	0.73	
T+3	0	0.53	
T+4	0	0.39	
T+5	0	0.29	
Mean			
(T to T+5)	0	0.01	

Table 4.	The	Growth	Rate of	of Consu	mption	Following	a Drop	in the (Growth I	Rate
of Incom	e									
Percent										

Source: Based on author's calculations as described in the text. The expected growth rate of income starts at 2 percent and then falls to 0 percent annually, starting at year T.

learns, before deciding his year T consumption, that the growth rate of his permanent income will henceforth be 0 percent rather than 2 percent. In year T, consumption will immediately drop because the present discounted value of future income is now lower. In years T + 1 and beyond, the growth rate of consumption will converge toward the growth rate of income. The table computes the growth rates of consumption that would occur if this consumer experienced the mean values of all shocks in years T through T + 5.

After the initial shock of -2.85 percent, the growth rate of consumption converges fairly rapidly to the growth rate of income. By T + 3, consumption growth is 0.53 percent and by T + 5, it is only 0.29 percent. Over the entire period from T through T + 5, the average growth rate is 0.01 percent—very close indeed to the 0 percent growth rate of income. From this evidence, the buffer-stock model is thus capable of quantitatively as well as qualitatively explaining the consumption/income parallel.

The CEQ-PIH Model and the Cyclical Behavior of the Personal Saving Rate

According to the simplest CEQ-PIH model, in periods when income is temporarily low or temporarily declining, the saving rate should also be low or temporarily declining. Recessions, of course, are periods of



Index, business cycle peak = 100

Figure 8. NIPA Personal Saving Rate around the Business Cycle Peak

Source: Author's calculations based on saving data from National Income and Product Accounts (NIPA) and NBER-dated business cycle peaks. Index at peak = 100. For each postwar recession, the personal saving rates for the four quarters preceding and following the NBER peak quarter are divided by the rate in the peak quarter. The figure presents the average of the resulting figures for each quarter between four quarters prior to and four quarters after the peak (multiplied by 100).

temporarily low or temporarily declining income, yet figure 8 shows that, on average, the personal saving rate as defined by the National Income and Product Accounts (NIPA) does not decline as the economy enters a recession.⁶⁰ A related point is made by figure 9, which shows the ratio of household debt to personal disposable income.⁶¹ If the simplest CEQ–PIH model were a complete explanation of consumer spending, the ratio of debt to consumption should be increasing during recessions because people should borrow to cover expenses while income is tem-

60. The NIPA personal saving rate includes expenditures on consumer durables as consumption. Economic theory, however, suggests that there is an important sense in which durables spending is more like saving than like consumption. Because durables spending is highly procyclical, if saving is redefined to include durables spending, saving will be procyclical too. This issue is discussed in greater detail below.

61. The general pattern of debt deceleration during recessions occurs in all major subcategories of debt: mortgage liabilities, consumer installment credit, other consumer credit, and other debt.


Source: Author's calculations based on Federal Reserve Board, *Balance Sheets for the U.S. Economy*, and NIPA. Figure plots quarterly data as a centered three quarter moving average. Vertical lines represent NBER-dated business cycle peaks.

porarily low.⁶² Instead, households typically *reduce* debt burdens during recessions.

Precautionary saving motives provide a natural explanation for why consumers do not reduce saving or increase borrowing during recessions. Anecdotal evidence of this abounds. A Gallup poll in Britain in August 1991, while Britain was in recession, found that 60 percent of households thought it was a good time to increase their savings, while 73 percent expected unemployment to increase. One of the poll directors said "consumers everywhere were inclined towards 'precautionary saving' to provide a cushion against the threat of unemployment."⁶³ A

62. Of course, the model's implications are really about net worth rather than debt. It is true that net worth is procyclical, but I concentrate on the behavior of debt because asset revaluations make interpretations of the cyclical behavior of household net worth difficult. It thus remains possible that the apparently anomalous behavior of debt is due to portfolio shifts and does not fundamentally constitute a violation of the CEQ-PIH model.

63. "Consumer Caution Constrains Recovery," *The Times of London*, August 7, 1991, business section, p. 20.

May 1992 editorial in the U.S. trade publication *Automotive News* corroborated such sentiments:⁶⁴

[F]olks still aren't buying cars. . . . And I am convinced that most Americans are still concerned about their jobs. As long as that insecurity exists, we are going to see a sluggish auto industry.

These quotations suggest that unemployment expectations have a direct effect on the saving rate. However, the CEQ-PIH model provides no direct role for unemployment expectations; only expected income matters. Formally, Campbell showed that under his model, current saving, S_t , should equal the present discounted value of expected declines in income:⁶⁵

(6)
$$S_t = -E \sum_{i=1}^{\infty} R^{-i} \Delta Y_{t+i},$$

where Y is income and R is the discount factor.

Campbell found some aggregate evidence supporting the proposition that consumers save more when future income growth is slower; he called this phenomenon "saving for a rainy day." Thus if expectations about future unemployment and future income growth could be measured directly, a valid test of the CEQ–PIH model would be to see whether the current saving rate depends on unemployment expectations after controlling for income expectations. I conducted such tests, first using data drawn from solely from the University of Michigan *Surveys* of *Consumers*,⁶⁶ and next combining the survey data with aggregate data on personal saving and unemployment. In both cases, unemployment data are shown to have an important influence on saving, even after controlling for income expectations.

Consumers participating in the *Surveys of Consumers* are asked whether over the next one or two years they expect their own household real income to increase or decrease, and whether over the next year they expect the unemployment rate to increase or decrease. They are also asked: "If there were a major purchase that you wanted to make, do you think that now is a time when it would be okay to use some of your sav-

64. Automotive News, May 11, 1992, p. 12.

66. University of Michigan, *Surveys of Consumers*. The data exist for every month since January 1989 and for periodic months prior to that date.

^{65.} Campbell (1987).

ings, or is now a time when you would be especially reluctant to use some of your savings?" I constructed three indexes: one for the fraction of consumers who believe it is "not okay" to draw down their savings (SAV); one for the fraction who believe the unemployment rate will increase (MU); and one for the fraction who believed household income would increase (MY).

These questions, of course, do not correspond directly to the theoretical constructs in Campbell's equation, but with some imagination it is possible to use them to perform a loose test of the model. It seems reasonable to interpret SAV_t as a proxy for saving S_t on the theory that people are "especially reluctant" to draw down their liquid assets during periods when they would be more likely to save in order to increase those liquid assets. More formally, SAV_t can be interpreted as an indicator of the shadow value of liquid assets. Under this interpretation, it should be positively correlated with saving in the form of liquid assets.

I will interpret MY as a proxy for the expected changes in income over the next year or two, $MY_t = \gamma E_t \sum_{i=1}^n \Delta Y_{t+i}$, where *n* is either 12 months or 24 months depending on how one wishes to interpret the question. In this case, because *R* is close to one, $MY_t \approx \gamma \sum_{i=1}^n R^{-i} \Delta Y_{t+i}$. Thus defining $\alpha = 1/\gamma$, we can rewrite Campbell's equation as

$$S_t \approx \alpha M Y_t + R^{-n} \alpha M Y_{t+n} + R^{-2n} \alpha M Y_{t+2n} \dots + e_t,$$

where e_t is the error term in time period t.

Define $PDVMY_t = R^{-n} MY_{t+n} + R^{-2n} MY_{t+2n} + \dots$ Then Campbell's equation becomes:

(7)
$$S_t \approx \alpha \left(MY_t + PDVMY_t \right) + e_t,$$

and the prediction of equation 6 is that α is negative. Of course, *PDVMY*_t is not observable, so we cannot estimate equation 7. What we can estimate is

(8)
$$SAV_t = \alpha_0 + \alpha_1 M Y_t + v_t.$$

 $PDVMY_t$ is an omitted variable in this equation. If $corr(MY_t, PDVMY_t) \neq 0$, the coefficient α_1 will be biased, so it might appear to be impossible to test the CEQ-PIH model by estimating equation 8. This is unduly pessimistic, however, because if $corr(MY_t, PDVMY_t) \ge 0$, then the theory still predicts that $\alpha_1 < 0$. The assumption that $corr(MY_t, PV_t) \ge 0$.

 $PDVMY_t \ge 0$ amounts to an assumption that expectations of income growth over the short term are not negatively correlated with expectations of growth over the long term. A formal test of this assumption would require estimating a time-series process for MY so that values of MY could be projected into the infinite future to construct PDVMY. For instance, suppose MY follows an AR(1) process, so that $MY_t = a + b$ MY_{t-1} . If b is positive, then corr $(MY_t, PDVMY_t) > 0$. In fact, when such equations are estimated, b is indeed positive.

A simpler approach, however, is to estimate $MY_t = a + b MY_{t-n}$, because in this case if $|R^{-n} b| < 1$, we have $PDVMY_t = k + [R^{-n} b/(1 - R^{-n} b)] MY_t$, where k is an uninteresting constant. If we estimate equation 8, the expected value of α_1 will be $\alpha_1 = \alpha [1 + (R^{-n} b/1 - R^{-n} b)]$. The CEQ-PIH model implies that $\alpha < 0$, so if b > 0, then the theory implies that $\alpha_1 < 0$. When this equation is estimated on monthly data, the result is

$$MY_t = 45.5 + 0.02 MY_{t-12}.$$
(8.2) (0.22)

Thus, with enough assumptions, it is possible to coax from the CEQ– PIH model the implication that $\alpha_1 < 0$. Regression 1 in table 5 presents the results of estimating equation 8, and, as predicted, α_1 is estimated to be negative and statistically significant; households express greater reluctance to dip into savings in periods when they are pessimistic about income growth.

Now consider estimating the following regression:

(9)
$$SAV_t = \alpha_0 + \alpha_1 M Y_t + \alpha_2 M U_t + v_t.$$

Unemployment expectations entered nowhere in equation 6, so according to the CEQ–PIH model, α_2 should be zero so long as corr(MU_t , $PDVMY_t$) = 0. In a buffer-stock model, by contrast, a higher probability of job loss increases the target buffer stock w^* , so heightened job insecurity should be associated with a higher personal saving rate in both the short and the long run.⁶⁷ As regression 2 of table 5 shows, when equation

^{67.} It is not clear, however, whether the degree of employment insecurity depends more on the level of the unemployment rate or on the expected change in the unemployment rate. For the moment, however, I am considering only variables in the Michigan survey, which means MU, the expected change in the unemployment rate. In analysis below I will also include the level of the unemployment rate.

Regres- sion ^a	Dependent variable ^b	Independent variables ^c		AR(1) serial		Box-	Number of
		MY	MU	correlation coefficient	$\overline{\mathbf{R}}^2$	Pierce p-value	observa- tions
(1)	SAV	-0.61 -(1.97)		•••	0.08	0.00	118
(2)	SAV	-0.27 -(1.36)	0.42 (5.88)		0.49	0.00	118
(3)	SAV	-0.20 -(1.37)	0.38 (4.12)	0.38 (2.42)	0.58	0.98	41
(4)	CRED	-0.32 -(1.16)			0.01	0.00	120
(5)	CRED	0.03 (0.15)	0.42 (6.18)		0.38	0.00	120
(6)	CRED	-0.08 -(0.54)	0.29 (4.08)	-0.24 (-0.14)	0.32	0.46	41

 Table 5. Regressions Relating Saving or Credit to Expected Changes in Unemployment

 and Income

Source: Author's calculations using data from the University of Michigan, Surveys of Consumers, various issues. Numbers in parentheses are *t*-statistics. The data are at a monthly frequency and different series are available over somewhat different periods. The sample period for all data ends in July 1992. All series are available every month beginning in January 1989, but before that time some series are available only in scattered months. The number of observations therefore varies from regression to regression because different methods of estimation impose different requirements for availability of lagged values. For instance, calculating a first difference requires observations in two adjacent months, while calculating an AR(1) restricts the sample to the period over which data were available in every month. A constant term was also estimated but is not reported.

a. Regressions 1, 2, 4, and 5 display evidence of serial correlation, so the reported *t*-statistics use autocorrelation and heteroskedasticity-consistent standard errors calculated using the procedure described in Newey and West (1987). This procedure is valid if the form of the autocorrelation is an MA process of limited length. I chose a lag of eight.

b. I created the dependent variables used in the regressions from answers to questions asked in the Michigan surveys. The exact questions asked in the surveys were as follows:

"If there were a major purchase that you wanted to make, do you think that now is a time when it would be okay to use some of your savings, or is now a time when you would be especially reluctant to use some of your savings?" SAV = (fraction answering not okay) minus (fraction answering okay).

SAV = (fraction answering not okay) minus (fraction answering okay). "If there were something that you wanted to buy, do you think that now is a time when it would be okay to buy on credit, or is now a time when you would be especially reluctant to take on new debt?"

CRED = (fraction answering not okay) minus (fraction answering okay).

c. The independent variables reflect answers to the following questions from the Michigan survey: "[In] the next year or two, do you expect that your (family) income will go up more than prices will go up, or less than prices will go up?"

MY = (fraction answering "more than prices") minus (fraction answering "less than prices").

"How about people out of work during the next 12 months—do you think that there will be more unemployment than now, about the same, or less?"

MU = (fraction answering "more" unemployment) minus (fraction answering "less" unemployment).

9 is estimated, α_2 is highly significant and positive, while α_1 , although still negative, is no longer statistically significant. The obvious interpretation is that the CEQ model is wrong because the expected change in the unemployment rate affects current saving, even after controlling for expectations about income growth. However, it remains possible that MU_t matters only because it is correlated with the omitted variable, $PDVMY_t$.

To test this possibility, I ran regressions of the form $MY_t = a + b$ $MY_{t-n} + c MU_{t-n}$. If the only reason MU is related to SAV is that MU is a good indicator of expectations about future changes in income, we should expect to see a strongly negative coefficient c on MU.⁶⁸ In fact, MU is only very weakly negatively related to future expected changes in income:

$$MY_t = 46.0 + 0.02 MY_{t-12} - 0.02 MU_{t-12}$$
(7.9) (0.18) -(0.33)

This regression suggests that it is unlikely that the reason MU is so strongly correlated with SAV is because MU is a good predictor of PDVMY. This does not amount, of course, to a rigorous test of whether regression 2 of table 5 is consistent with the CEQ model; such a test would be valuable, but is beyond the scope of this paper.

The Box-Pierce statistic for regression 2 indicates substantial evidence of serial correlation, so regression 3 reports the results when AR(1) serial correlation in the errors is assumed. The AR(1) coefficient is highly significant, but there is no evidence of further serial correlation. The coefficient on the expected change in income is diminished further, and is now well below the threshold for statistical significance. The unemployment expectations variable, however, remains overwhelmingly statistically significant.

In modern econometrics, it is often considered inappropriate to address a problem of serial correlation by simply assuming an AR(1) error process because the model being tested often does not imply that serial correlation in the error term should exist. Indeed, the CEQ–PIH model does not provide any explicit justification for serial correlation in the saving rate. However, in a buffer-stock model, the saving rate *should* be highly serially correlated because the personal saving rate is a function of the net wealth ratio, an omitted variable that certainly should be highly serially correlated. In this sense, the existence of strong serial correlation might be construed as further evidence in favor of the bufferstock model.

68. In this case, under further simplifying assumptions, $PDVMY_i = K + [R^{-n}b/(1 - R^{-n}b)MY_{t-n}] + R^{-n}c/(1 - R^{-n}b)MU_{t-n}$. When equation 9 is estimated, the resulting estimates of α_1 and α_2 should be $\alpha_1 = \alpha[1 + (R^{-n}b/1 - R^{-n}b)]$, and $\alpha_2 = \alpha (R^{-n}c/(1 - R^{-n}b))$.

Economists are skeptical of polls, both because it is difficult to know whether respondents understand the questions as economists would, and because often what people say is not what they do. It is important, therefore, to know whether data on actual spending corroborate the results of table 5. Table 6 therefore presents regressions in which the dependent variable is the aggregate personal saving rate, s, rather than SAV.⁶⁹ Regression 1 reproduces the result of regression 1 in table 5; saving is negatively associated with expected income growth, as predicted by the CEQ-PIH model. However, unlike table 5, the coefficient is not significant here. Regression 2 adds two variables: MU, the expected change in the unemployment rate, and U, the current level of the unemployment rate. Both the level and the expected change in unemployment are included because it is not clear which is likely to be a better indicator of households' fears about job loss. My own intuition is that job fears are probably greater when the unemployment rate is unchanging at 10 percent than when the unemployment rate is 5 percent but increasing; this would imply that U_t might matter more than MU_t . Regression 2 of table 6 finds that both are highly statistically significant with positive signs: saving is higher both when the current level of the unemployment rate is high and when the unemployment rate is expected to rise.

Regression 2 is subject to the same critique leveled at regression 2 in table 5: it is possible that U_t and MU_t are significant only because they are correlated with $PDVMY_t$, rather than because expectations about unemployment have a direct effect on saving. As above, a crude test of this is to run a regression of the form $MY_t = a + b MY_{t-n} + c MU_{t-n} + d U_{t-n}$. If c and d were strongly negative, it would add credibility to the hypothesis that MU_t and U_t explain current saving only because they provide a signal about income after the period covered by MY_t . The results of such a regression are⁷⁰

$$MY_{t} = 45.7 + 0.09 MY_{t-4} - 0.05 MU_{t-4} - 34.8 U_{t-4}.$$

(4.7) (0.58) - (0.14) - (0.62)

69. This is calculated as the difference between total disposable personal income and total personal consumption expenditures—both defined as in Blinder and Deaton (1985)— divided by total disposable labor income. This is the measure used by Campbell and Deaton (1989).

70. In this regression, n = 4 rather than 12 as above because these regressions use quarterly rather than monthly data. The *t*-statistics are corrected for serial correlation using the Newey-West procedure with eight lags.

Regressions ^b	Dependent variable	МҮ	MU	U	AR(1) serial correlation coefficient	Time trend	$\overline{\mathbf{R}}^2$	Box- Pierce p value	Number of observations
(1)	S	-6.19 -(1.09)				-7.71 -(2.76)	0.37	0.00	62
(2)	S	8.75 (2.08)	2.78 (2.15)	1.10 (7.32)		-4.98 -(4.63)	0.65	0.00	62
(3)	S	0.95 (0.28)	0.89 (0.68)	0.67 (2.73)	0.61 (5.34)	-7.28 -(3.68)	0.78	0.59	61
(4)	S	-13.96 -(2.50)	• • •	• • •		-11.10 -(5.67)	0.55	0.00	58
(5)	S	13.75 (1.82)	-0.24 -(0.41)	0.91 (3.21)		-3.17 -(3.17)	0.64	0.00	57
(6)	S	-1.42 -(0.18)	-1.47 -(0.63)	0.15 (0.38)	0.67 (3.28)	-10.37 -(3.16)	0.78	0.49	56

Table 6. Regressions of the Personal Saving Rate on Income Expectations, Unemployment Expectations, and the Unemployment Rate^a

Source: Author's calculations using data from University of Michigan, Surveys of Consumers, various issues, NIPA, and Bureau of Labor Statistics, Employment and Earnings. The dependent variable for all six regressions is the personal saving rate, s. The numbers in parentheses are t-statistics. A constant term was also estimated, but is not reported.

a. Definitions for aggregate variables follow Blinder and Deaton (1985). U is the unemployment rate, s is the personal saving rate calculated by subtracting total personal consumption expenditures from total disposable personal income, dividing by disposable labor income, and multiplying by 100 to convert to percentage points. MU is unemployment expectations and MYis income expectations as defined in table 5; both variables are converted to a quarterly basis. The sample period, determined by data availability, is 1977:1-1992:2. The Michigan variables were reported on a monthly basis; they were converted to quarterly data by taking the average of monthly readings within each quarter. If there were no observations within a quarter (as was the case for a few quarters of M), the previous quarter's observation was repeated.

Using these definitions of the data, the table would be read as follows. Regression 2, for example, would be interpreted as saying that if MU moved from 0 (the number of consumers expecting the unemployment rate to rise equals the number expecting it to fall) to 1 (all respondents say the unemployment rate will rise), the personal saving ln rate would rise by 2.78 percentage points. If the current unemployment rate were to rise by 1 percentage point, the personal saving rate would rise by 1.10 percentage points.

b. Regressions 4, 5, and 6 are calculated using instrumental variables. The instruments include one full year (lags 2 through 5) of MY, MU, U, $\Delta \ln Y$, and yield, where $\Delta \ln Y$ is the change in log personal disposable income, and "yield" is defined as the difference between the return on three-month and ten-year T-bills. The yield variable was inspired by the work of Stock and Watson (1991), who found it to be a useful leading indicator. Regressions 1, 2, 4, and 5 display evidence of serial correlation, so the reported *t*-statistics use autocorrelation and heteroskedasticity consistent standard errors. As in table 5, I chose eight lags. The reported standard errors did not appear to be very sensitive to the choice of lag length. The results are similar to those presented above for MU using the monthly data: MU_t and U_t are negatively associated with future MY_t , but the relationship is quite weak, so it appears unlikely that MU and U are useful in explaining the current saving rate only because of their ability to forecast $PDVMY_t$. Again, it would be valuable to test this hypothesis more rigorously, but such an exercise is beyond the scope of this paper.

As with regression 2 of table 5, the residuals of regression 2 in table 6 exhibit strong evidence of serial correlation.⁷¹ Regression 3 therefore estimates a similar equation allowing for an AR(1) error term in the equation. The coefficient on U_t remains positive and strongly significant, but the significance of MU_t is diminished substantially.

Another possible objection to the results of tables 5 and 6 is that they may stem from reverse causality: perhaps exogenous increases in the saving rate *cause* the expected subsequent increases in the unemployment rate. To see how this could happen, consider an IS/LM economy in which consumers are rational but subject to random shocks to taste that move consumption. A negative consumption shock means high saving now, but the drop in C reduces aggregate demand. As demand falls, firms fire some workers and pay their remaining workers less. So an increase in saving is followed by drops in employment and income.

Regressions 4, 5, and 6 address this problem by reestimating regressions 1, 2 and 3 using instrumental variables. Using variables dated time t - 2 and earlier as instruments removes the simultaneity problem, as long as the shock to taste is not correlated with variables dated t - 2 or earlier. For regression 5, the results are similar to those of regression 2: the predicted unemployment rate is strongly positively associated with the current saving rate. For regression 6, however, which allows for AR(1) serial correlation in the error term, none of the three explanatory variables is statistically significant. This means that it is not possible to rule out reverse causality as an explanation for the dependence of the personal saving rate on current and expected future unemployment.

Another criticism of these regression results might be in the definition of saving as the difference between disposable labor income and total personal consumption expenditures, which includes expenditures for

71. However, an augmented Dickey-Fuller test is able to reject the hypothesis that the errors from this equation have a unit root.

durable goods. The theory of durable goods implies that purchases of durables might be more appropriately thought of as saving than consumption, because a durable good provides a store of value (an "asset") that is "consumed" over many subsequent periods. Thus much empirical macroeconomic work has concentrated on consumption of nondurables and services, or the ratio of nondurables and services consumption to income.⁷² If the fraction of consumption expenditures devoted to nondurables and services consumption were stationary, then it might be possible to construct a "nondurables and services" saving ratio, as Campbell attempts to do.⁷³ However, this ratio appears to follow a random walk with drift, so I was unable to construct a sensible indicator of the saving rate using expenditures on nondurables and services.

It is possible, therefore, that the results of table 6 are caused solely by a sharp decline in purchases of durable goods when job fears worsen. If durables purchases were a perfect substitute for other forms of saving, the assertion that saving increases when job fears worsen would then be unjustified; the correct statement would be that the mix of saving shifts away from acquisition of durable goods to other forms of saving. I should note that this is not an implication of the CEQ–PIH model of durables as developed, for instance, by Mankiw.⁷⁴ However, the fact that substantial transactions costs are associated with the purchase and sale of durable goods means that they are not a perfect substitute for other assets. They are a particularly bad substitute for the kind of liquid assets that are most useful as a buffer stock against uncertainty. As a first approximation, therefore, I would argue that for the purposes of testing for the effects of employment uncertainty on saving, the NIPA definition of saving used in table 6 is probably the appropriate definition.

From the evidence of tables 5 and 6 it appears that, even controlling for expectations about future income growth, consumers both express a greater desire to save and actually save more when they believe that the unemployment rate will be rising. They also save more when the level of the unemployment rate is high. Although these results cannot defini-

72. See, for example, Campbell (1987) and Cochrane (1991).

73. Campbell (1987).

74. Mankiw (1981). It is, however, an implication of the (S, s) model of durables as developed in Eberly (1992). She finds that that model implies that spending should be highly sensitive to uncertainty. Her empirical tests confirm this prediction, although using measures of uncertainty that do not directly reflect job fears as here.

Brookings Papers on Economic Activity, 2:1992

tively rule out either the CEQ-PIH model or an IS/LM model in which rational consumers are subject to shocks in taste, it seems fair to say that the most straightforward interpretation is that people save more when fears about job security increase—a natural implication of the bufferstock model, but not of the CEQ-PIH model or the Keynesian model.

Unemployment Expectations and the Campbell-Mankiw Model

Even if the results of tables 5 and 6 were entirely inconsistent with the CEQ-PIH framework, this might not be very interesting because a great many papers have already rejected a wide variety of implications of that model (although rarely, as here, in favor of an alternative model in which consumers are in a sense *more* rational and sophisticated than in the CEQ model). It would therefore be interesting to test a model that has not already been rejected many times. The Campbell-Mankiw model is the obvious candidate. Campbell and Mankiw state: "We expect that the simple model presented here . . . will be hard to beat as a description of the aggregate data on consumption, income, and interest rates."⁷⁵

Regressions 1 and 3 of table 7 reproduce the basic result of Campbell and Mankiw for both total consumption and consumption of nondurables and services: spending responds significantly to predictable changes in income, even using twice-lagged instruments. Regressions 2 and 4, however, show that when the instrumented value of the University of Michigan *Surveys of Consumers* unemployment expectations variable is added to the right hand side of the equation, it is highly statistically significant; moreover, the change in income is less significant and has a smaller coefficient than before the expectations variable was added.

The unemployment expectations variable has a negative coefficient, implying that consumption growth is slower in periods when consumers are pessimistic about future unemployment conditions. Intuition might indicate that this is favorable evidence for the idea that precautionary saving is important. Recall, however, that the Euler equation for consumption under uncertainty was

(5)
$$\Delta \ln C_{t+1} = \rho^{-1} (r - \delta) + \frac{1}{2} E_t \operatorname{var} (\Delta \ln C_{t+1}) + e_{t+1}.$$

75. Campbell and Mankiw (1989, p. 212).

	Dependent		Indep vari	oendent ables ^c		Box-	
Regressiona	variable ^b	Constant	$\Delta \ln Y$	MU	$\overline{\mathbf{R}}^2$	p-value	
(1)	$\Delta \ln C$	0.195 (2.54)	0.717 (6.88)		-0.02	0.08	
(2)	$\Delta \ln C$	0.507 (4.95)	0.403 (3.41)	-1.273 -(3.65)	0.35	0.01	
(3)	$\Delta \ln CNS$	0.525 (9.65)	0.484 (6.41)		0.08	0.88	
(4)	$\Delta \ln CNS$	0.743 (10.47)	0.267 (2.99)	-0.887 -(3.50)	0.39	0.58	
(5)	$\Delta \ln CN$	0.341 (2.76)	0.013 (0.10)	-0.906 -(1.91)	0.04	0.02	
(6)	$\Delta \ln CFOOD$	0.222 (1.85)	0.114 (0.63)	-1.008 -(2.17)	0.09	0.45	
(7)	$\Delta \ln CREC$	0.960 (4.27)	0.098 (0.47)	-0.793 -(1.06)	-0.02	0.00	

Table 7.	Regressions of	of the	Change in	Consumption	on tl	he Change	in	Income
and the	Unemploymen	t Exp	ectations I	ndex				

Source: Author's calculations using data from the University of Michigan, Surveys of Consumers, and NIPA. Numbers in parentheses are t-statistics.

a. Regressions are instrumented using the same set of twice-lagged instruments described in table 6. As in tables 5 and 6, the r-statistics were calculated with heteroskedasticity and autocorrelation-consistent standard errors using the Newey-West procedure with eight lags and a damping factor of one. All variables are in 1987 dollars. The dependent variable and the Δ In Y variable were multiplied by 100 so that the regression results could be interpreted as percentage points per quarter.

b. The dependent variables are defined as follows:

C is total real personal consumption expenditures following Blinder and Deaton (1985),

CNS is total real PCE for nondurable goods and service, excluding clothing, following Blinder and Deaton (1985), who argue that at the quarterly frequency, clothing is a durable good,

CN is total real PCE for nondurable goods, excluding clothing,

CFOOD is total real PCE for food, and

CREC is total real PCE for personal recreation services.

c. The independent variables are defined as follows:

 $\Delta \ln Y$ is the change in log of personal disposable income, and

MU is unemployment expectations, as defined in table 5.

According to this equation, consumption should grow *faster* in periods when uncertainty (the expected variance of future consumption growth) is higher. This counterintuitive result is explained by figure 6, presented in the third section. The reaction to an increase in uncertainty is for the saving rate to increase as soon as uncertainty increases (which was confirmed by the results in tables 5 and 6). Thereafter consumption should grow slightly faster than before as the wealth stock is built up toward its target level. Thus, in periods when uncertainty is high *but unchanging*, consumption growth should be high (at least until the new target wealth stock is achieved).





Source: Author's calculations. The model consumption rate depicts an optimal target path in which the consumer earns expected income for the first five periods and unemployment expectations change in the sixth period, as in figure 6. The moving average consumption path depicts a moving average of current consumption and two lagged values of the optimal consumption path.

A speculative explanation of these results is given in figure 10. If it is difficult to adjust consumption downward instantaneously, the level of consumption might adjust slowly to an increase in uncertainty. The line labeled "model consumption path" simply reproduces the consumption path calculated in figure 6 for the optimal behavior following an increase in the expected probability of unemployment. The line labeled "moving average consumption path" shows the average of the current value and the last two values of the "model consumption path." The moving average consumption path is designed to capture the idea that consumption might adjust slowly, rather than instantly, to the new unemployment expectations. In the moving average consumption path, consumption falls both in the initial period and in the two following periods (when no change in expectations occurred). Other than its predictions for the growth rate of consumption immediately following a shock, the qualitative implications of this model are largely similar to those of the model with instantaneous adjustment. However, formal tests of the buffer-

stock model would find that consumption displays "excess smoothness" with respect to current changes in unemployment expectations, and could also find that the predictable level of unemployment is important for current consumption growth.

Certainly slow adjustment must be the rule for some important categories of nondurables and services consumption. For instance, one of the largest components of services consumption is imputed rent on owner-occupied housing. Imputed rent changes either when a house is sold or when it is revalued. Given the transactions costs associated with buying and selling houses, it would be unreasonable to expect this component of consumption to adjust downward instantly, or within a single quarter, upon news of a higher probability of future unemployment. Similarly, expenditures on household utilities are probably determined largely by the characteristics of the home and might be difficult to adjust quickly. Many other components of consumption also plausibly may be difficult to adjust instantly because of precommitment or fixed costs.

Campbell and Deaton and many others have found that consumption appears to exhibit "excess smoothness" with respect to changes in permanent income.⁷⁶ Figure 10 suggests that the results of table 7 could be reconciled with the results of tables 5 and 6 if there is also "excess smoothness" of consumption with respect to changes in unemployment expectations. It would not be surprising if a common explanation for both kinds of excess smoothness existed. One possibility is given by a model of Ricardo Caballero.⁷⁷ He interprets it as a model of near-rationality, but it could just as easily be interpreted as one in which even nondurable consumption spending has fixed adjustment costs. He finds that such a model can generate excess smoothness of consumption with respect to shocks to permanent income. Such a model would also generate excess smoothness with respect to changes in uncertainty, thus providing at least a potential for a joint explanation of both kinds of excess smoothness.

Another possible explanation is habit formation in consumption, as suggested by James Dusenberry, John Heaton, George Constantinides, and others.⁷⁸ Casual observation convinces me that it is not uncommon even for freshly minted economics Ph.D.'s to be slow to adjust their con-

^{76.} Campbell and Deaton (1989).

^{77.} Caballero (1992).

^{78.} See, for example, Dusenberry (1949), Heaton (1990), and Constantinides (1990).

sumption to a level commensurate with their income after beginning a professional career.

A few pieces of evidence on these topics are given by the last three regressions of table 7. The idea here is to choose categories of consumption that a priori seem less subject to precommitment or fixed costs of adjustment, and to see whether the MU term is less significant for such components. Among the broad aggregates of consumption, nondurables spending is probably easier to adjust than services consumption; regression 5 shows that the expected change in unemployment is less significant in explaining nondurables consumption than for nondurables and services combined. Disaggregating further within nondurable goods, food seems a good candidate for a category of spending that can be adjusted quickly; indeed, regression 6 shows that the coefficient on unemployment expectations is about as significant in explaining food expenditures as in explaining total nondurable spending. Within the services category, "recreational services" includes movie and theater tickets and a host of other categories of spending that, plausibly, are discretionary; regression 7 shows that there is no significant evidence that recreational services grow more slowly when unemployment expectations are predictably high.

This is tenuous evidence. However, in combination with the earlier evidence that the saving rate rises when unemployment expectations increase, it suggests that the reaction of most categories of nondurables and services consumption to an increase in uncertainty is roughly as shown in the "moving average consumption path" in figure 9: slower consumption growth and higher saving for at least a few quarters.

The tests of this section are particularly interesting because of the nature of the evidence they provide against other models. In the past ten years, most rejections of the permanent-income hypothesis using macroeconomic data have been interpreted as resulting from a failure of consumers to be sufficiently forward-looking, or from the existence of binding liquidity constraints. The results of this section, however, suggest that consumers are both forward-looking and able to adjust their saving in response to perceptions about the future, but that they are forwardlooking in a different (and, for that matter, more sophisticated) way than postulated by simple certainty-equivalence models.

The fundamental message of these results and those of the previous section is that precautionary saving is important in understanding the

cyclical behavior of consumption and saving. Models that do not allow separate effects for the future level of income and for the degree of uncertainty about future income (particularly uncertainty about employment) seem to miss important features of the data.

Applying the Results

This final section considers whether the buffer-stock model can shed some light on two case studies that represent puzzles for more traditional models of consumption and saving. The first puzzle is the muchanalyzed decline in the personal saving rate over the past fifteen to twenty years. I argue that two important culprits may be the decline in the growth rate of personal income and the relaxation of borrowing standards that began in the late 1970s. The second puzzle is why consumption growth has been so weak very recently. Using models similar to those estimated in the previous section, I argue that continuing (and justified) pessimism about unemployment accounts for a substantial part of the weakness.

The Secular Decline in the Personal Saving Rate

Figure 11 displays U.S. personal saving as a fraction of personal disposable income since 1960. A host of studies has tried to explain the secular decline in the saving rate that began sometime in the mid-1970s. A literature spawned by Martin Feldstein emphasizes the importance of the increased generosity of the Social Security system, which reduces the need to save for retirement.⁷⁹ Others have cited demographic effects of the baby boom generation,⁸⁰ the changing behavior of pension plans,⁸¹ the effect of increases in household net worth,⁸² reduced bequest saving by the elderly,⁸³ and a host of other factors. No consensus has emerged.

In the third section, I showed that the personal saving rate in a bufferstock model will tend toward $s^* \approx gw^*$. The buffer-stock model there-

- 80. Bovenberg and Evans (1989).
- 81. Bernheim and Shoven (1988).
- 82. Wilcox (1991).
- 83. Auerbach, Kotlikoff, and Weil (1992).

^{79.} Feldstein (1974).



Figure 11. U.S. Personal Saving Rate, 1960:1–1992:2 Personal saving as a percent of disposable income

fore suggests two natural categories of explanations for the long-term decline in s: a fall in g or a fall in w^* . I will first consider whether the falling personal income growth rate can explain the drop in the personal saving rate, and, next, whether various trends in the economy may have reduced w^* . Leading candidates are a reduction in household income uncertainty because of the increasing prevalence of two-earner households, and a relaxation of borrowing constraints following financial deregulation in the late 1970s and 1980s.

FALLING INCOME GROWTH AND THE SAVING RATE. From 1960 to 1980, the annual growth rate of per capita personal income in the United States was about 2.5 percent, yet from 1980 to 1991, the growth rate was only about half as large, at 1.3 percent. The average personal saving rate was 7.3 percent in the earlier period and 6.0 percent in the latter period. Although low growth is often blamed on low saving, it is possible in a buffer-stock model for the causality to go in the opposite direction: slow growth can cause low saving. This is a surprising result because in a traditional LC/PIH model, lower income growth implies lower lifetime income, lower consumption, and hence *higher* saving rates.

Source: Author's calculations using NIPA. NIPA definitions of saving were used. Vertical lines represent NBERdated business cycle peaks. Figure plots quarterly data as a centered three-quarter moving average.

First I must show that it is theoretically possible for lower growth to lead to a lower saving rate. Formally, the question is whether the derivative of the target personal saving rate with respect to g is positive or negative:

(10)
$$\frac{ds^*}{dg} \approx \frac{d}{dg}gw^* = w^* + g\frac{dw^*}{dg}.$$

At g = 0, the second term on the right hand side vanishes. So starting at an income growth rate of zero, it is clear that if w^* is positive, an increase in the growth rate of income must increase the personal saving rate. However, there is no guarantee that the relationship between saving and growth is positive for all values of g because the optimal buffer stock is smaller when income growth is faster. (See the discussion of figure 4 and table 3 in the third section.) Thus dw^*/dg is negative.

Away from g = 0, therefore, the sign of equation 10 is theoretically ambiguous and must be determined by numerical solution of the model. Referring again to table 3, we can calculate that the target personal saving rate gw^* is zero at g = 0; 0.02 times $0.44 \approx 0.9$ percent at g = 0.02; and 0.04 times $0.37 \approx 1.5$ percent at g = 0.04. Thus under the parameter values chosen here, the relationship between income growth and the target saving rate is positive between g = 0 and g = 0.04. However, these same calculations show that the model is not quantitatively capable of explaining the entire decline in the U.S. personal saving rate. The difference in the saving rate at g = 0 and at g = 2 percent was only 0.9 percent. So a change in income growth of 2 percent, almost twice as large as the actual 1.2 percent change, generates a change in the saving rate of only 0.9 percent, slightly smaller than the actual change of 1.3 percent. Under the base parameter values, therefore, the buffer-stock model implies that the drop in income growth can at most account for a bit less than half the decline in the personal saving rate in the 1980s.

The basic reason that the model cannot explain the entire drop in saving is that the target net wealth ratio w^* is not large enough. If in reality consumers are more risk averse, or have a lower discount rate, or for any other reason have a higher w^* , then the model would be able to explain more of the decline in saving as the result of the decline in growth.⁸⁴

^{84.} I have concentrated here on long-run comparisons of saving rates across target ratios, rather than the short-term dynamics of the saving rate, because the short-term dynamics go the wrong way: the short-term reaction to a decline in the growth rate of income is to *increase* saving.

EVIDENCE FOR REDUCED TARGET WEALTH. Before turning to the evidence on factors that may have reduced the target net wealth ratio, w^* , it may be useful to review the short-term and long-term effects of such a reduction. Suppose that in period T a change in the economic environment occurs (say, a sudden and permanent loosening of liquidity constraints) that results in a lower target for net wealth; after period T, no further changes in the environment occur. Relative to their new target, consumers now hold too much net wealth. The result could colorfully be described as a consumption binge: there will be a sudden sharp increase in consumption and a sharp decline in the saving rate as consumers begin to spend down their net wealth toward the new target. As the binge period wears off (and net wealth declines), the saving rate will gradually recover somewhat, but its new target gw^* will be lower than before the drop in w^* .

The analysis is somewhat more difficult in the case of gradual changes in the economic environment that slowly reduce w^* over time. For instance, imagine a gradual decline in income uncertainty. The easiest way to think of this is to imagine consumers going on a small consumption binge every year as w^* continues to drop slightly every year. Two features of such a scenario are clear: first, the personal saving rate should drop over time as w^* drops, and, second, if the decline in w^* stops, the saving rate should subsequently rise somewhat as consumers end their consumption binges. Again, however, the eventual target saving rate gw^* will be lower than it was before the decline in w^* began.

Turning now to factors that may have reduced the target wealth ratio w^* over the last twenty years, two broad categories of explanations seem worth investigating: declines in income uncertainty and relaxation of liquidity constraints. I will treat these in order.

Declining income uncertainty. One of the most important changes in the nature of household earnings process in the past thirty years has been the remarkable increase in labor force participation by women. The fraction of married women who are in the labor force increased from 30.5 percent in 1960 to 40.8 percent in 1970 to 57.8 percent in 1989.⁸⁵ The household income of a typical (two wage-earner) household today is therefore probably less uncertain than that of a typical (one

85. Statistical Abstract of the United States, 1991 (table 643, col. 1, p. 391).

wage-earner) household twenty or thirty years ago because it is unlikely that both husband and wife would lose their jobs at the same time. Eventually, of course, this trend must end, because female labor force participation cannot exceed 100 percent. It is not clear, however, whether the trend has ended yet.

A second development is the expansion of insurance. (Insurance reduces the need for precautionary saving by reducing uncertainty.) The buffer-stock model would suggest that unemployment insurance may be particularly important; in fact, the unemployment insurance program was expanded considerably during the 1970s. In 1972 about 18 percent of workers were not covered by unemployment insurance; since reforms in the late 1970s, all but about 3 percent of the workforce has been covered.⁸⁶ The model would predict that as the program became more comprehensive, the need for buffer wealth should have declined, and thus the personal saving rate should have declined. The saving rate did indeed decline fairly sharply in the late 1970s, although the timing of the decline did not coincide exactly with changes in unemployment insurance. The model would also predict that as the changes in the program came to an end, the saving rate should rise, as it did in the early 1980s. The final prediction would be that the average level of the saving rate in the new equilibrium should be lower than before the expansion in the system. Indeed the average personal saving rate since the early 1980s has been substantially lower than previously.

Despite the apparent match between the model and the facts, it strains credibility to attribute too much of the pattern of personal saving to the described changes in the unemployment insurance system. Furthermore, during the 1980s the unemployment insurance system became somewhat less generous as eligibility requirements were tightened in many states.⁸⁷ If unemployment insurance were the vital force driving the personal saving rate, at least some increase in the saving rate should have occurred following these changes. However, the saving rate in the late 1980s continued to decline. A balanced view might be that the increase in unemployment insurance is probably responsible for a modest part of the decline in saving, but by no means most of the decline.

86. U.S. Department of Labor (1972, p. 8) and Burtless (1983, p. 229).

87. For instance, the number of months workers are required to have been employed in order to qualify for benefits was increased.

Another exceptionally important form of insurance is health insurance, which has also expanded greatly since the 1950s and 1960s. The fraction of the population covered by the most common form of health insurance, hospital insurance, rose from 68 percent in 1960 to 82 percent in 1980, although it had fallen back to 76 percent by 1985.⁸⁸ The fraction of the average medical bill that was paid directly by the patient (as opposed to being paid by insurance) fell from 55 percent in 1960 to 28 percent in 1985.⁸⁹ Thus the expansion of health insurance could have contributed to the long-term downward trend in the personal saving rate. Again, however, during the middle and late 1980s, the fraction of the population covered by health insurance declined and health insurance benefits became less generous, so trends in health insurance are not much help in explaining why saving was particularly low in the 1980s.

A more speculative possibility is that the long expansion of the 1980s and the long decline in the unemployment rate led consumers to revise doward their estimated probability of unemployment or the estimated variance of income shocks. Certainly the depth and duration of the recent recession came as a surprise to consumers and economists alike. If consumers become permanently more pessimistic about the unemployment rate, this could lead to some recovery in the saving rate in both the short run and the long run as w^* is revised up. The next section of the paper considers whether this phenomenon might be partly responsible for the exceptional weakness of consumption growth recently.

Relaxation of borrowing constraints. Although the buffer-stock model described formally above does not incorporate liquidity constraints, the last part of the third section argued that a model with explicit liquidity constraints could exhibit behavior very similar to that described for the model presented here. Fundamentally, I view the model with explicit liquidity constraints as an alternative and slightly more complex manifestation of the buffer-stock model, rather than as a separate and competing model. The most important reason to model liquidity constraints are relaxed. But, as I argued above, the basic qualita-

^{88.} Health Insurance Association of America (1987, p. 10).

^{89.} See Summers and Carroll (1987, p. 628).

		··· <u></u>			
Time period	Total debt	Mortgage debt	Installment credit	Other consumer credit	Other debt
1961-65	0.68	0.45	0.14	0.05	0.04
1966-70	0.71	0.45	0.16	0.05	0.05
1971-75	0.70	0.44	0.17	0.04	0.05
1976-80	0.73	0.48	0.17	0.03	0.05
1981-85	0.76	0.51	0.16	0.03	0.06
1986–90	0.89	0.62	0.20	0.02	0.05

 Table 8. Types of Debt Relative to Disposable Income, 1961–90

 Ratio

Source: Author's calculations using data from Federal Reserve Board, Balance Sheets for the U.S. Economy, and NIPA. Other debt is total debt minus mortgage debt, installment credit, and consumer credit.

tive result is clear: when liquidity constraints are relaxed, impatient buffer-stock consumers will borrow more and will have a lower target wealth. Relying on this admittedly rather loose description of the liquidity constrained version of the model, I will proceed to argue that part of the decline in the personal saving rate since the late 1970s has been caused by a relaxation of borrowing constraints and a consequent increase in debt.

During the 1980s debt certainly did increase. Table 8 provides details of the changes by debt category since the 1960s. The ratio of total liabilities to income crept up by about 0.3 percentage points a year between 1961–65 and 1976–80. That ratio grew by about 1.5 percentage points a year between 1976–80 and 1986–90.

This rapid growth of debt concided with a comprehensive liberalization and deregulation of financial markets. Casual evidence suggests that credit became easier to get and that repayment terms became more lenient. For example, the fraction of households owning at least one bank credit card increased from 16 percent in 1970 to 54 percent in 1989.⁹⁰ The average length of new car loans increased steadily, from 37.6 months in 1975 to 54.6 months in 1990.⁹¹ The average down payment as a fraction of house value for first-time home buyers declined from 18 percent in 1976 to 11.4 percent in 1985.⁹²

90. Canner and Luckett (1992, p. 656).

91. Federal Reserve Statistical Release, G19, "Consumer Installment Credit," various releases.

92. Statistical Abstract of the United States (1987, table 1293, p. 716).





Percent of outstanding debt that is deliquent

Source: Author's calculations based on the Mortgage Bankers Association, National Delinquency Survey, various years. Vertical lines are NBER-dated business cycle peaks.

The natural inference is that liberalization is at least partly responsible for the increase in debt. However, debt growth could instead reflect demographic factors, such as the beginning of home purchases by the baby boom generation, or portfolio shifts in which households chose to hold both more assets and more debt, or a number of other factors. The best proof that consumers who were poorer credit risks have had increased access to credit is therefore evidence about the incidence of delinquency, default, or other nonpayment of debt. Figure 12 shows the fraction of mortgages whose payments were delinquent by 90 days or more since 1955. This figure trended steadily upward from the early 1970s to the mid-1980s.⁹³ Figure 12 tells a similar story for delinquencies on auto loans and shows higher delinquency rates for bank credit cards

93. The peak around 1985 occurred because of the regional recession in oil-producing regions of the country and thus may not be directly related to a trend toward higher average default rates. However, even after allowance is made for this factor, a clear upward trend in delinquencies appears between the early 1970s and the mid- to late 1980s.

in the mid-1980s and early 1990s than at any other time in the data sample. Data on other delinquency rates, foreclosure rates, personal bankruptcies, and other measures confirm the impression that in the mid- to late 1980s, debt repayment problems increased substantially.

Financial intermediaries may have had several reasons to increase their willingness to lend to questionable borrowers in the 1980s. The first is the perverse incentives that eventually led to the savings and loan crisis: because deposits were insured, owners of savings and loan institutions (S&L's) had every incentive to make risky loans. Broadly speaking, if the loans paid off, the S&L owners profited; if the loans failed, the government took the loss. It is clear that S&L's made very risky loans to commercial borrowers. S&L's also substantially expanded their market presence in unsecured consumer lending in the mid-1980s. Mortgages, however, were the bread-and-butter business of S&L's, and through the mid- to late 1980s, S&L's continued to originate approximately 60 percent of home mortgages. Even a modest relaxation of lending standards by S&L's could therefore have made a big difference in this market. On the other hand, the evidence is not clear whether the new S&L's mortgage loans from this period were any riskier than mortgage loans made by banks in the same period. Most of the riskiest lending by S&L's was to the commercial sector, not to homeowners.

Another important development was the expansion of the mortgagebacked securities markets facilitated by government agencies such as the Federal Home Loan Mortgage Corporation, the Federal National Mortgage Association, and the Government National Mortgage Association. The increased liquidity of mortgage loans provided by these agencies apparently allowed banks to relax credit standards for loans they originated. The banks could then bundle the loans together and sell them to investors, thus avoiding some of the risk associated with mortgage lending.

Several other factors may have played a role. One important development was that improvements in information technology made it easier to swap information on prospective borrowers. Marco Pagano and Tullio Jappelli report that the number of credit reports grew from 60 million in 1960 to 100 million in 1970 to 400 million in 1990.⁹⁴ Standard

^{94.} Pagano and Jappelli (1991, p. 27).

models of credit rationing, such as that of Joseph Stiglitz and Andrew Weiss, rely on the lender's inability to distinguish between "good" and "bad" borrowers.⁹⁵ Such models can generate outcomes in which the lender will not lend at all if he cannot distinguish between good and bad borrowers. However, if he can distinguish between borrowers, he will lend to the good borrowers at low interest rates and the bad borrowers at high interest rates. Therefore, if the improvement of credit reporting technologies made distinguishing between good and bad credit risks easier, credit should have become more accessible to everyone. Such information technology developments probably were one of the factors behind the deluge of credit card applications that began appearing with great regularity in many people's mailboxes in the mid- to late 1980s.

Whatever the reason for looser credit, in an economy populated by buffer-stock consumers who face borrowing constraints, loosening those constraints should reduce the saving rate in both the short run and the long run. Easier credit may be the most plausible explanation for the low saving rate of the 1980s. If this explanation is correct, the saving rate should recover somewhat (although not to its original level) when credit stops loosening. In fact, many of the trends that apparently led to the loosening of credit during the 1980s had either halted or begun to reverse by late 1989 or 1990, around the beginning of the "credit crunch." Perhaps not coincidentally, the personal saving rate has been gradually improving since that time, averaging 4 percent in 1989, 4.3 percent in 1990, 4.7 percent in 1991, and 5.2 percent in the first half of 1992. However, interpretation of recent movements in the saving rate is complicated by the fact that the economy has been in recession.

Precautionary Saving and the Recent Recession

The flip side of the increasing saving rate recently has been slow consumption growth. In particular, consumption growth has been remarkably lackluster since the trough⁹⁶ of the recent recession.⁹⁷ A common

95. Stiglitz and Weiss (1981).

96. Although the NBER dating committee has not yet met at the time of this writing, I will assume that the trough will eventually be dated as the second quarter of 1991.

97. This claim is supported by the fact that macroeconomic models and economic forecasters have consistently overpredicted consumption growth over this period.

theme of press discussion of the weakness in consumption has been that "debt overhang" from the 1980s has been restraining consumer spending. The discussion just concluded confirmed that in the 1980s the ratio of consumer debt to personal consumption did rise at an unprecedented rate. Furthermore, the rapid growth in the debt ratio ended near the beginning of 1990, and the economy slipped into recession soon after.

What often remains unexamined is why a level of debt that was acceptable in 1989 and was freely chosen by consumers at that time constituted "debt overhang" in 1990 or 1991. The natural explanation in the buffer-stock framework is that the target level of net worth has increased, making consumers uncomfortable with levels of wealth that seemed adequate in happier times. In particular, the two characteristic features of recessions—a low income growth rate g and a higher unemployment rate—should both increase to the target level of wealth. Which feature is more important will depend on the parameters of the model, but the simulations above showed that under the parameter values used here, the target wealth ratio was considerably more sensitive to the probability of unemployment than to the growth rate of income.

The regressions of tables 5 and 6 suggest that both effects occur: saving is higher when future income growth is low and when prospects for employment are low. However, those regressions have direct implications only for net worth, and not explicitly for debt. The problem is that changes in debt could just represent portfolio shifts if, for instance, most changes in debt happened because consumers paid off debts by drawing down assets, rather than saving more.

To gain some insight on consumer attitudes toward debt, the University of Michigan *Surveys of Consumers* ask another interesting question: "If there were something that you wanted to buy, do you think that now is a time when it would be okay to buy it on credit, or is now a time when you would be especially reluctant to take on new debt?"⁹⁸ This question is similar to the one on saving reported above, but the answers are at least potentially different. In fact, regressions 4 through 6 of table 5 show that attitudes toward debt are broadly similar to attitudes toward saving: consumers are reluctant to increase debt when they expect unemployment to rise. However, the relationship between expected income growth and attitudes toward debt is much weaker.

^{98.} University of Michigan, Surveys of Consumers, various years.



Source: Author's calculations based on University of Michigan, Surveys of Consumers (1992). Index is equal to the percent of consumers who expect unemployment to rise minus the percent who expect unemployment to fall, multiplied by 100. See table 5 for a further explanation. Vertical lines represent NBER-dated business cycle peaks and troughs. The 1991 trough was estimated by the author as 1991:2.

This analysis suggests that, to the extent that a "debt overhang" problem has existed, the problem is not so much the debt itself as that consumers' expectations have changed in ways that made them uncomfortable holding the debt that they had voluntarily assumed previously. The regressions of tables 5 and 6 strongly suggest looking at unemployment expectations. Thus figure 13 plots the unemployment expectations index over time, along with NBER business cycle peaks and troughs. Although the index did not display especially unusual behavior in the beginning of the most recent recession, it has failed to drop steeply since the trough, as it usually has after previous recessions. This is strongly suggestive, because forecasts of consumption did not begin to go seriously off track until the trough.⁹⁹

99. For the MPS model, for example, consumption growth began to go off track around the beginning of the recession, but the model began making large errors only after the trough.

122

Figure 13. Expected Unemployment Index

Forecast type ^b	Dependent variable	Model	Average residuals since peak, 1990:3–1992:2	Average residuals since trough, 1991:2–1992:2
Out-of-sample	Δc	Campbell-Mankiw	-0.57	-0.37
Out-of-sample	Δc	Augmented	-0.48	-0.25
Out-of-sample	Δcns	Campbell-Mankiw	-0.52	-0.42
Out-of-sample	Δcns	Augmented	-0.46	-0.34
Out-of-sample	Δcd	Campbell-Mankiw	-0.86	-0.02
Out-of-sample	Δcd	Augmented	-0.62	0.29
In-sample	Δc	Campbell-Mankiw	-0.49	-0.49
In-sample	Δc	Augmented	-0.29	-0.22
In-sample	Δcns	Campbell-Mankiw	-0.46	-0.48
In-sample	Δcns	Augmented	-0.31	-0.28
In-sample	Δcd	Campbell-Mankiw	-0.68	-0.49
In-sample	Δcd	Augmented	-0.22	0.14

Table 9.	Average	Forecast	Errors	for the	Campb	ell-Mankiw	Model	and the	Augmented
Campbel	l-Mankiw	v Model s	ince the	Last B	usiness	Cvcle Peak	a		

Source: Author's calculations using data from the University of Michigan, *Surveys of Consumers*, and NIPA. Variables are defined as in table 7, except for *cd*, which is defined as log total real expenditure on consumer durables (including clothing).

a. This table is read as follows. The first line indicates that the Campbell-Mankiw model (Campbell and Mankiw, 1989), whose parameters were estimated on data from 1961:1 through 1990:2, overpredicted consumption growth by an average of 0.57 percent per quarter in the period from 1990:3 to 1992:2, and by 0.37 percent per quarter over the period from 1991:2 to 1992:2. The next line indicates that the augmented model overpredicted on average by 0.48 percent and 0.25 percent, respectively, over the same period.

b. "Out-of-sample" designates forecasts made with model parameter values estimated on the period before the peak of the latest business cycle. "In-sample" designates forecasts made with model parameters estimated using the entire sample through 1992:2.

To look more rigorously at the role of pessimism about employment conditions in the recent weakness of consumption growth, I reestimated the models reported in table 7 for total consumption and nondurables and services consumption; I then estimated a similar model for durables spending. I estimated the models through the second quarter of 1990 (the quarter before the peak) and then used them to predict consumption on a rolling basis through the second quarter of 1992. (The estimation results were very similar to those reported in table 7, so I did not report them.) The first panel of table 9 shows that the straight Campbell-Mankiw model overpredicted consumption growth by an average of 0.57 percent per quarter since the peak in 1990:3, and by 0.37 percent per quarter since the trough of the recession in 1991:2. By comparison, the augmented model, which includes the unemployment expectations variable, overpredicted by an average of 0.48 percent per quarter since the peak, but only 0.25 percent per quarter since the trough. Thus, in a crude

sense, it might be fair to say that consumer pessimism about unemployment accounts for about one-third of the error in the forecast of growth in total consumption spending of the standard Campbell-Mankiw model.

The next part of table 9 shows how the two models performed for nondurables and services spending and for durables spending. The augmented model performed slightly better in predicting nondurables and services consumption, and performed far better in predicting durable goods spending. Indeed, while the standard model underpredicts durables spending substantially, the augmented model *overpredicts* durables spending. This suggests that, controlling for consumers' pessimism about unemployment, durables spending has actually been stronger than would have been anticipated.

The next part of table 9 performs a similar analysis, but uses versions of the models estimated through the end of the sample in 1992:2. As should be expected, the in-sample forecasts are somewhat better than the out-of-sample forecasts. For total consumption expenditures, the augmented model's average error in the period since the peak is less than half the Campbell-Mankiw model's average error.

It is tempting to draw the conclusion that lingering consumer pessimism is the cause of the recent weakness of the economy. However, the pessimism has turned out to be justified: unemployment remains stubbornly high and has even increased substantially since the trough. There are two possible extreme interpretations of these events. The first is that consumer expectations are largely determined by random shocks, but that expectations turn out to be correct because they constitute self-fulfilling prophecies. Pessimism leads to low spending, which in turn generates layoffs. Thus in the end, the pessimism turns out to have been justified. The second extreme possibility is that the future unemployment rate is autonomous and unaffected by current consumption shocks. If consumers are rational and forward-looking, however, current changes in consumption will anticipate subsequent changes in unemployment without being the underlying cause of those changes, in the same way that one's decision about whether to bring an umbrella to work may be a good forecast of whether it rains, but certainly does not cause the rain. Although there are probably elements of truth in both interpretations, it is difficult to know which is a better description of the recent behavior of the economy.

A small clue may have been provided by the events following the end

of the Persian Gulf War in the spring of 1991, when a surge of consumer confidence and consumer spending occurred for several months. If ever there were a natural experiment to determine whether exogenous shocks to consumption can generate subsequent economic growth sufficient to sustain and justify the original consumption shock, this was it. But by the fall of 1991, it was clear that the spending of the previous spring had not generated enough new employment and income to sustain a normal recovery. The economy began to struggle again, and consumer confidence collapsed. Although not conclusive, this experience suggests that the recent weakness in consumption growth is largely a rational reaction to justified pessimism about continuing high unemployment.

Directions for Future Research

Uncertainty about future income has been largely left out of much macroeconomic analysis of consumption and saving, in part because no convenient model existed that could analyze uncertainty. This paper shows that the buffer-stock model can be used effectively to explain a variety of macroeconomic phenomena, both at the cyclical frequency and over longer time spans. The model implies in particular that the unemployment rate, or expectations about the future unemployment rate, may be quite important in determining current consumption and saving behavior, because unemployment fears are likely to be the most important component of uncertainty in overall household income. The empirical results confirm that consumers' self-reported expectations about the unemployment rate are closely associated with the level of saving and the growth rate of consumption. Although the correlations between consumption growth and unemployment expectations do not conform to the strict interpretation of the model, if consumption does not adjust instantaneously to changing unemployment expectations, even these correlations can be explained as roughly consistent with the model.

The model also provides several explanations for the drop in the personal saving rate in the 1980s: lower income growth, greater income security, and easier borrowing. Of these, easier borrowing seems likely to have been the most important. Finally, the model can explain why over time horizons of a few years or more, consumption growth and income growth move so closely together. If consumers have a target net wealthto-income ratio, in the long run, consumption growth must match income growth; otherwise the net wealth-to-income ratio would explode or implode.

The buffer-stock model does not, of course, explain all behavior of all consumers. Some consumers accumulate very large amounts of wealth; neither the buffer-stock model nor any traditional consumption-smoothing model is likely to explain the behavior of these consumers. Many other consumers do engage in some form of life-cycle saving behavior, particularly as they approach retirement. Recall, however, that the required condition for buffer-stock saving in the infinite horizon buffer-stock model was $\rho^{-1} (r - \delta) < g$. Although in a finite horizon model the condition is more complex, this equation provides the correct intuition: whether buffer-stock saving behavior will occur depends on the expected future growth rate of income. If expected income growth is high early in life but lower (or negative) as retirement approaches, it is entirely possible that consumers will engage in buffer-stock saving when young but, after a certain age, will switch to more traditional life-cycle saving behavior as their expected future income growth falls.

The buffer-stock model presented here calls for a variety of further work, both theoretical and empirical. On the theoretical side, it would be desirable to develop the model with a choice of assets. For instance, consumers are able to make long-term but illiquid investments with a high rate of return such as housing. I would speculate that some set of parameter values can accommodate such behavior, predicting longterm investment in housing alongside buffer-stock behavior with respect to short-term fluctuations in income. Whether those parameter values would be plausible is hard to know in advance. Another interesting theoretical question is whether formally adding habit formation or fixed costs of adjustment to consumption can explain the fact that the average growth rate of consumption is negatively related to the level of unemployment expectations. Again, the answer is likely to depend on parameter values.

On the empirical side, a great deal more work could be done to flesh out the time-series relationship between saving, consumption, and unemployment expectations. This may be a difficult task because of the reverse causality problem, but that has not prevented a substantial research program on consumption and income expectations in the last few

years. A careful analysis of the University of Michigan *Surveys of Consumers* data on income expectations and unemployment expectations would also be interesting. It would be useful to know whether these series satisfy constraints that represent a rational expectations formation process. A substantial opportunity also exists for research using microeconomic data. In a 1992 paper, Andrew Samwick and I, for instance, find that households that face greater income uncertainty hold more assets, and in approximately the ways predicted by the buffer-stock model.¹⁰⁰ Another type of useful microeconomic research would be to provide a more detailed and careful description than provided in the second section about the stochastic process for household income, particularly around episodes of unemployment. In sum, a host of interesting questions remain about both the theoretical and empirical implications of the model.

APPENDIX

A. Estimating the Variance of Transitory and Permanent Shocks When the Transitory Shock Is Not Independently Identically Distributed

The first section of the paper assumes that the level of the transitory component of income is identically independently distributed (iid). The assumption that the *level* is iid implies that the *first difference* of log income should follow an MA(1) serial correlation process (because the transitory change in income this period should reverse itself next period). The estimates of John Abowd and David Card¹⁰¹ and Thomas Ma-Curdy¹⁰² are consistent with an MA(1) serial correlation process in the first difference, but both papers prefer an MA(2) to an MA(1) serial correlation. Hence, the assumption of an uncorrelated transitory component in levels is probably not strictly correct.

Fortunately, the decomposition methodology described in the text can be easily extended to the case of a transitory error that is MA(q)

- 100. Carroll and Samwick (1992).
- 101. Abowd and Card (1989).
- 102. MaCurdy (1981).

serially correlated rather than iid.¹⁰³ Suppose $\ln V_t$ is MA(q). For notational convenience call $r_m = \ln YL_{t+m} - \ln YL_t$. Then var $(r_m) = m \sigma_{\ln N}^2 + 2 \sigma_{\ln V}^2$ still holds for any m > q because cov $(\ln V_{t+m}, \ln V_t) = 0$. So long as the data contain at least q + 2 years of data, it will still be possible to solve two linear equations for $\sigma_{\ln N}^2$ and $\sigma_{\ln V}^2$.

If there are $\tau + q$ years of data, there are $(\tau - 1)$ potential *m*'s and $(\tau - 2)$ potential different estimates of $\sigma_{\ln N}^2$ and $\sigma_{\ln V}^2$. Assuming that the serial correlation properties of income are constant through time, each individual estimate is consistent, but greater efficiency is obtained by averaging the separate estimates of $\sigma_{\ln V}^2$ and $\sigma_{\ln N}^2$. Suppose q = 2 and $\tau = 8$. Call the vector of calculated variances $v = [var(r_3), \ldots, var(r_{10})]'$. In a linear regression of v on a constant and the vector (3, 4, ... 10)', the coefficient on the constant term should be $2\sigma_{\ln V}^2$ and the coefficient on the constant term should be $2\sigma_{\ln V}^2$ and the methodology actually used for estimating $\sigma_{\ln V}^2$ and $\sigma_{\ln N}^2$.

B. Deriving the Optimal Policy Rule for Consumption

The model presented in the third section is solved using backwards recursion on the period-by-period Euler equations to derive a numerical optimal consumption rule for periods $T-1, T-2, \ldots, T-n$.¹⁰⁵ Optimal consumption in any period will depend only on the sum of current assets and current income, a variable that I will call X, gross wealth,

$$X_t = W_t + YL_t,$$

where W_t is net wealth and YL_t is noncapital labor income.

The evolution of gross wealth over time is given by

$$X_{t+1} = R(X_t - C_t) + YL_{t+1},$$

where C_t is consumption at time t and R is the interest factor.

Now consider a consumer solving the finite horizon problem, with last period of life T. Because the consumer is assumed to spend what-

104. As in constructing the *YLRATIO*, the predictable life-cycle element of changes in income was removed before these calculations were performed.

105. The solution method follows that of Deaton (1991).

^{103.} An AR process would be more difficult to address. However, Abowd and Card, MaCurdy, and most other economists have consistently rejected an AR specification in favor of MA specifications.

ever gross wealth he has in the last period, the Euler equation for optimal consumption in the period before last is

(B-1)
$$u'(C_{T-1}) = R\beta E_{T-1}u'(C_T) = R\beta E_{T-1}u'(X_T),$$

where β is the discount factor.

This equation is in the levels of C, X, and YL. The income process is given by

$$YL_t = P_t V_t$$
$$P_{t+1} = G P_t N_{t+1},$$

where G is the growth factor for permanent income; N_t is the transitory shock; P_t is permanent labor income; and V_t is a multiplicative transitory shock in year t.

It would be more convenient to express equation B-1 in terms of ratios to permanent income P. To accomplish this, it is necessary to write out the difference equation for x_i , the ratio of X_i to permanent income P_i , as a function of c_i , the ratio of C_i to permanent income P_i :

(B-2)

$$\begin{aligned} x_{t+1} &= X_{t+1}/P_{t+1} \\ &= [R(X_t - C_t) + YL_{t+1}]/P_{t+1} \\ &= [R(x_t - c_t) P_t + YL_{t+1}]/P_{t+1} \\ &= [R(x_t - c_t) P_t]/G P_t N_{t+1} + YL_{t+1}/P_{t+1} \\ x_{t+1} &= R(x_t - c_t)/G N_{t+1} + V_{t+1}. \end{aligned}$$

This means that B-1 can be rewritten as

(B-3)
$$u'(c_{T-1}P_{T-1}) = R\beta E_{T-1}u'(x_TP_T) = R\beta E_{T-1}u'(x_TGP_{T-1}N_T).$$

For a homogeneous utility function, it is permissible to divide both arguments by P_{T-1} . Doing this and substituting for the value of x_T yields

$$u'(c_{T-1}) = R\beta E_{T-1} u' \{ [R(x_{T-1} - c_{T-1})/GN_T + V_T] GN_T \}.$$

For any given value of the gross wealth ratio x_{T-1} , this equation implicitly defines the optimal value of the consumption ratio $c_{T-1}(x_{T-1})$. Now between period T - 1 and T - 2 a similar Euler equation holds:

$$u'(C_{T-2}) = R\beta E_{T-2} u'(C_{T-1})$$

$$u'(c_{T-2} P_{T-2}) = R\beta E_{T-2} u'[c_{T-1}(x_{T-1}) P_{T-1}],$$

$$u'(c_{T-2}) = R\beta E_{T-2} u'\{c_{T-1}[R(x_{T-2} - c_{T-2})/GN_{T-1} + V_{T-1}]GN_{T-1}\}.$$

So this implicitly defines c_{T-2} as a function of x_{T-2} , c_{T-2} (x_{T-2}). More generally,

(B-4)
$$u'(c_t) = R\beta E_t u' \{c_{t+1}[R(x_t - c_t)/G N_{t+1} + V_{t+1}] G N_{t+1}\},\ u'(c_t) = R\beta \int_N \int_V u' \{c_{t+1}[R(x_t - c_t)/G N_{t+1} + V_{t+1}] G N_{t+1}\} dF(V) dF(N).$$

The function $c_t(x_t)$ is defined implicitly by this equation. Because there is no analytical solution for $c_t(x_t)$, numerical methods must be used to solve; for details of how the numerical calculations are performed, see my 1992 paper.¹⁰⁶ Thus the optimal consumption rules are found by solving successively for $c_{T-1}(x_{T-1})$, $c_{T-2}(x_{T-2})$, and so on back to the first period of life. As illustrated in figure 3, however, under the parameter values chosen here, the c_{T-i} functions converge rather quickly; the function they converge to is designated c(x).

C. Proving That Consumption Growth Is Higher When Wealth Is Lower

If shocks to consumption are lognormally distributed, the Euler equation for consumption growth is given by:

(C-1)
$$\Delta \ln C_{t+1} \approx \rho^{-1} (r - \delta) + \frac{1}{2} \rho E_t \operatorname{var} (\Delta \ln C_{t+1}) + e_{t+1},$$

where C_t is consumption; ρ is the coefficient of relative risk aversion; r is the interest rate; and δ is the discount rate.

In the text, I asserted that the expected growth rate of consumption will be greater when wealth is lower because the conditional variance of consumption growth is higher when there is less wealth available to buffer income shocks. I will prove this for the simplified case where the interest rate is zero and income shocks are independently identically distributed, $\tilde{Y} = \bar{Y} + u$, where \bar{Y} is mean income and u is an iid mean zero error (there is no income growth). In this case, there will be an optimal consumption rule C(X) that relates the level of consumption to the level of gross wealth X. Define

(C-2)
$$Z(X) = \ln C(X)$$

$$\sigma_{\Delta z}^{2}(X_{t}) = E_{t} [\operatorname{var} (\Delta \ln C_{t+1})] = E_{t} \{\operatorname{var} [Z(\tilde{X}_{t+1}) - Z(X_{t})] \}$$

$$= E_{t} \{\operatorname{var} [Z(\tilde{X}_{t+1})] \},$$

106. Carroll (1992).

where \tilde{X}_{t+1} is the random variable representing the value of gross wealth in period t + 1, and the last equality is true because X_t is known at time t so that var $[Z(X_t)] = 0$. Then the task is to prove that $\frac{d}{dx_t} \sigma_{\Delta z}^2(X_t) < 0$, the expected variance of consumption growth declines as gross wealth increases. Define $F(X_t) = E_t \tilde{X}_{t+1} = X_t - C(X_t) + \overline{Y}$. Then $Z(\tilde{X}_{t+1}) =$ $Z[F(X_t) + u_{t+1}]$. Taking a first order Taylor expansion yields

(C-3)
$$Z(\tilde{X}_{t-1}) \approx Z[F(X_t)] + Z'[F(X_t)] u_{t+1}$$

But because $Z(F[X_t])$ is known at time t and has no variance, it follows that

(C-4)
$$E_t[\operatorname{var} [Z(\tilde{X}_{t+1})]] \approx E_t(\operatorname{var} \{Z'[F(X_t)] | u_{t+1}\}) = \{Z'[F(X_t)]\}^2 \sigma_u^2$$

Thus the sign of $\frac{d}{dx_t} \sigma_{\Delta z}^2(X_t)$ will be the same as the sign of $\frac{d}{dx_t} \{Z'[F(X_t)]\}^2$.

(C-5)
$$\frac{d}{dx_t} \{ Z'[F(X_t)] \}^2 = 2 Z'[F(X_t)] \frac{d}{dx_t} \{ Z'[F(X_t)] \}$$

Now $Z(X) = \ln C(X)$, so $Z'(X) = \operatorname{so} Z'(X) = C'(X)/C(X)$, which will always be positive because both the level of consumption and the marginal propensity to consume (MPC) out of wealth are always positive. The question is now to determine the sign of

(C-6)
$$\frac{d}{dx_t} \{ Z'[F(X_t)] \} = Z''[F(X_t)] F'(X_t).$$

Taking the derivative of Z'(X) = C'(X)/C(X) and dropping arguments for clarity gives

(C-7)
$$Z'' = \frac{C''C - C'C'}{C^2}.$$

Kimball's work¹⁰⁷ established that for utility functions that exhibit *decreasing absolute prudence* (as does the CRRA utility function used by Zeldes¹⁰⁸ and the model in this paper) C''(X) < 0; that is, the marginal propensity to consume declines as gross wealth increases. However, the

107. Kimball (1990a,b).

108. Zeldes (1989a).
level of consumption C is always positive, so Z'' must be always negative. The remaining question is the sign of $F'(X_t)$.

(C-8)
$$\frac{d}{dx_t}F(X_t) = \frac{d}{dx_t}[X_t - C(x_t) + \overline{Y}] = 1 - C'(X_t).$$

Because the MPC out of wealth will always be less than one, $F'(X_t) > 0$. Thus C-6 is negative, implying C-5 is negative, implying that d/dx_t $\sigma_{\Delta z}^2 < 0$. Thus, wealthy consumers have a lower variance of consumption changes, and so a lower expected growth rate of income.

D. Proof That If $\mathbf{x}_t > \mathbf{x}^*$, Then $\mathbf{E}_t \mathbf{x}_{t+1} < x_t$

In the text, I claim that if gross wealth x is greater than the target value, then the expected value of gross wealth next period will be less than gross wealth this period; formally, if $x_t > x^*$, then $E_t x_{t+1} < x_t$. This section proves that proposition for the case with only transitory shocks to income. Equation B-2 above established that

$$x_{t+1} = R[x_t - c_t(x_t)]/GN_{t+1} + V_{t+1},$$

where N is the shock to permanent income and V is the shock to transitory income. If there are no shocks to permanent income, $N \equiv 1$. If $H(x_t) = E_t x_{t+1}$, then, because V, by definition, has a mean of one, it is clear that

$$H(x_t) = R/G[x_t - c(x_t)] + 1.$$

Because x^* was defined as the point where $E_t x_{t+1} = x_t$, $H(x^*) = x^*$. Now consider some x_t in the vicinity of x^* , $x_t = x^* + u_t$. Taking a first order Taylor expansion,

$$H(x_t) \approx H(x^*) + H'(x^*) u_t = x^* + H'(x^*) u_t$$

So the question becomes whether $E_t x_{t+1} = x^* + H'(x^*) u_t < x^* + u_t = x_t$ —that is, whether $H'(x^*) < 1$. But

$$\frac{d}{dx_t}H(x_t) = R/G\left[1 - c'(x_t)\right].$$

Kimball's work¹⁰⁹ proved that the marginal propensity to consume is greater under uncertainty than under certainty. Using the usual approxi-

109. Kimball (1990a,b).

132

mations, the marginal propensity to consume under certainty is given by $[r - \rho^{-1}(r - \delta)]$, so that:

$$H'(x) = R/G [1 - c'(x_t)] < R/G [1 - r + \rho^{-1}(r - \delta)]$$

 $\approx (1 + r - g) [1 - r + \rho^{-1}(r - \delta)].$

So the question is now whether $(1 + r - g) [1 - r + \rho^{-1} (r - \delta)] < 1$. This condition will hold as long as $\rho^{-1} (r - \delta) < g$. But because ρ is positive and $(r - \delta)$ is negative for all parameter combinations considered in the paper, the left-hand side of the inequality is always negative, while g is always zero or positive. So, assuming all the approximations hold,

$$H'(x_t) < (1 + r - g) [1 - r + \rho^{-1} (r - \delta)] < 1,$$

which is the required condition to guarantee that if $x_t > x^*$, $E_t x_{t+1} < x_t$.

E. Proof That If $x_t = x^*$, Then $E_t \Delta C_{t+1} \approx g$

In the text, I claim that if gross wealth x_i is at the target ratio x^* , then the expected rate of growth of consumption is approximately equal to g, the growth rate of income. I prove this, again, for the simplified case where there are no shocks to permanent income.

Recall that (again using C for the level of consumption and c for the ratio of consumption to permanent income) $C_t = c_t P_t$, where P is permanent income. If there are no shocks to permanent income P, it grows so that $P_{t+1} = G P_t$. If $\ln (G) = g$, then the expected growth rate of consumption is given by

(E-1)
$$E_t [\ln (C_{t+1}) - \ln (C_t)] = E_t \{\ln [c(x_{t+1})P_{t+1}] - \ln [c(x_t)P_t]\}$$

 $= E_t [\ln c(x_{t+1})] + \ln (G) + \ln (P_t)$
 $- \ln c(x_t) - \ln (P_t),$
(E-2) $= E_t [\ln c(x_{t+1})] + g - \ln c(x_t).$

Recall from section B of the appendix that $x_{t+1} = R/G[x_t - c(x_t)] + V_{t+1}$. Recall that the mean of V was 1, and define $v_{t+1} = V_{t+1} - 1$. Call $F(x_t) = E_t x_{t+1} = R/G[x_t - c(x_t)] + 1$. Then define $Z(x) = \ln c(x)$ so that

(E-3)
$$\ln [c(x_{t+1})] = Z(x_{t+1}) = Z[F(x_t) + v_{t+1}].$$

A first-order Taylor expansion of Z around F(x) yields

(E-4)
$$Z(x_{t+1}) \approx Z[F(x_t)] + Z'[F(x_t)] v_{t+1}.$$

Brookings Papers on Economic Activity, 2:1992

The expected value of $Z(x_{t+1})$ yields $E_t Z(x_{t+1}) \approx Z[F(x_t)]$. Substituting this for the first term in equation E-2 yields

(E-5)
$$E_t [\ln (C_{t+1}) - \ln (C_t)] = E_t Z[F(x_t)] + g - \ln c(x_t),$$

$$\approx \ln c(E_t x_{t+1}) + g - \ln c(x_t).$$

But at the target ratio, $E_t x_{t+1} = x_t$, so the first and third terms are equal, so

(E-6)
$$E_t [\ln (C_{t+1}) - \ln (C_t)] \approx g$$

That is, at the target gross wealth ratio, the expected growth rate of consumption is approximately the same as the growth rate of income.

F. Proof That the Target Saving Ratio $s^* \approx gw^*$

In the text, I asserted that if wealth is equal to its target w^* , the personal saving rate will be approximately gw^* . I will prove this for the case in which there is no uncertainty in permanent income, $P_t = G P_{t-1}$ or $P_{t-1} = P_t/G$. Personal saving is defined as the difference between total personal income (including both capital and labor income) and consumption. Capital income in period t will be given by the interest rate r multiplied by the unconsumed portion of total period t - 1 resources, $(W_{t-1} + YL_{t-1} - C_{t-1})$. This leads to the definition of saving:

$$S_t = r (W_{t-1} + YL_{t-1} - C_{t-1}) + YL_t - C_t.$$

Because by definition

$$W_t = R(W_{t-1} + YL_{t-1} - C_{t-1}),$$

it follows that $S_t = (r/R) W_t + YL_t - C_t$.¹¹⁰

The precise question is the following. Suppose $w_t = w^*$. What is the expected personal saving rate in period t, where expectations are taken before the realization of the period t shock (that is, expectations are taken as of time t-1)?¹¹¹

110. A more familiar version of this equation would be $S_t = rW_t + YL_t - C_t$. The equation here differs because, in my notation, wealth in period t includes the interest income earned on savings from period t - 1. I made this notational choice because it usually simplifies the algebra of the model, but in this particular case it leads to an equation with a slightly nonstandard appearance.

111. Note that w_t is known at time t - 1, because $w_t = R/G(w_{t-1} + V_{t-1} - c_{t-1})$, so it makes sense to consider a case where w_t is known to equal w^* , even though the values of other time t variables is not yet known.

134

(F-1)
$$E_{t-1} s_{t} = E_{t-1} \left(\frac{r}{R} W_{t} + YL_{t} - C_{t}\right) / E_{t-1} YL_{t}$$
$$= E_{t-1} \left(\frac{r}{R} w_{t} P_{t} + V_{t} P_{t} - c_{t} P_{t}\right) / E_{t-1} V_{t} P_{t}$$
$$= \frac{r}{R} w_{t} + 1 - E_{t-1} c_{t}$$
(F-2)
$$E_{t-1} s_{t} \approx r(1-r) w_{t} + 1 - E_{t-1} c_{t}.$$

Now, noting that $x_t = w_t + V_t$, from equation B-2, it follows that

$$w_{t+1} = R/G (w_t + V_t - c_t)$$

$$E_{t-1} w_{t+1} = R/G (w_t + 1 - E_{t-1} c_t).$$

Now w^* was defined so that at $w_t = w^*$, it will be the case that E_{t-1} $w_{t+1} = w_t = w^*$. Substituting, it follows that:

$$w^* = R/G (w^* + 1 - E_{t-1} c_t)$$

$$1 - E_{t-1} c_t = w^* (G/R - 1)$$

(F-3)
$$1 - E_{t-1} c_t \approx w^* (1 + g - r - rg - 1) = w^* (g - r - rg).$$

Substituting this expression into equation F-2 at $w_{t-1} = w^*$ yields

$$E_{t-1} s^* \approx r (1 - r) w^* + (g - r - rg) w^*$$

= $rw^* - r^2 w^* + gw^* - rw^* - rgw^*$
= $(r - r^2 + g - r - rg) w^* \approx gw^*$,

where the last near-equality holds because r and g are both close to zero so that $r^2 \approx rg \approx 0$.

Comments and Discussion

Robert E. Hall: In this paper, Carroll provides an impressive review of the evidence on consumption within the framework of optimization theory. A basic point of this and his earlier work is that optimization theory can explain features of consumption behavior that previously seemed to require theories of nonoptimal behavior. Rather than invoke inefficient liquidity constraints, Carroll argues, it is better to look closely at optimal behavior for individuals who are trying to avoid the catastrophe of closeto-zero consumption. Further, Carroll states, what may have appeared to be spontaneous shifts of consumption and saving may actually be the rational response by consumers to changes in their stochastic environments.

An interesting feature of the paper is that it builds upon the same model of consumption that has dominated thinking in this area since the rational expectations revolution. Rather than advocate a new model, Carroll asks us to take seriously the nonlinearities in the existing model. He sees the majority of consumers as accumulating a buffer stock of net worth to guard against the small probability of earnings dropping to a level that would support only extremely low consumption. Once the buffer stock is in place, consumption has the same expected growth rate as income, contrary to, say, the certainty equivalence model. In the latter, only the present discounted value of earnings over the entire lifetime matters for consumption; spending does not respond at all to year-toyear changes in income, unless they signal changes in the entire present discounted value.

Households will not be in the buffer-stock mode unless they meet the following conditions. First, they must abhor low levels of consumption: utility must go to minus infinity as consumption goes to zero. The constant-elasticity family, including log utility, has this property, but not Figure 1. Supply and Demand for Capital



quadratic utility. Second, the rate of impatience must be high relative to the real interest rate or the rate of growth of income must be high. The precise condition is that the rate of impatience must be greater than the real interest rate less the product of the rate of growth of income and the coefficient of relative risk aversion. When this condition is satisfied, the household with quadratic preferences would go substantially into debt. By contrast, the household that puts high value on avoiding very low levels of consumption would not go into debt.

The difference between the two types of preferences in general equilibrium is shown in figure 1, which portrays the supply and demand for capital as functions of the real interest rate. The demand schedule is the marginal product of capital, net of depreciation. The supply schedule is the level of wealth that consumers will choose to hold in the steady state at a given interest rate. The upward-sloping curve describes the behavior of families with the constant relative risk aversion (CRRA) preferences hypothesized by Carroll. As the real interest rate approaches the critical value, the wealth held by consumers approaches infinity. But even at low interest rates, families hold positive amounts of wealth because of the buffer-stock motive. The horizontal dashed line is the wealth supply curve for families with quadratic preferences. Their supply is perfectly elastic at the critical interest rate. They would go indefinitely into debt if the interest rate were lower than the critical level.

From the figure, it is immediately apparent that the buffer-stock economy will have a higher capital stock and a lower real interest rate than the certainty-equivalence economy.

Carroll builds the model in both of the ways prevalent in the modern literature on consumption. First is the structural consumption function, which shows how families choose consumption given its determinants. Second is the Euler equation, which relates next year's expected marginal utility of consumption to this year's, based on the principle that the expected marginal rate of substitution should equal the expected price ratio. Carroll's structural consumption function must be calculated numerically to do full justice to the nonlinearities stressed in his approach. The calculations confirm that consumption and saving depend on net worth in such a way that there is a target ratio of net worth to permanent income.

Carroll's Euler equation (equation 5) shows that consumption growth depends on the two factors stressed in the Euler branch of the literature—the real interest rate, and the unpredictable random increment embodying the effect on consumption of all new information—plus a factor implicit in earlier work. The third factor is the conditional variance of next year's consumption, given information available this year. Variations over time in the conditional variance have been an important topic in the applications of the Euler equation in finance.¹ Carroll appears to be the pioneer in making the conditional variance an active part of the Euler equation in the study of consumption.

Carroll observes that an increase in uncertainty will trigger an immediate drop in consumption. In the structural consumption function, the amount of uncertainty is one of the determinants of the function; more uncertainty brings a downward shift. In the Euler equation, more uncertainty shows up as a negative value of the random disturbance term. Following the immediate drop, consumption growth rises; consumption eventually returns to its previous level, with consumers holding more wealth. In the structural consumption function, it is rising wealth that

^{1.} See Backus and Gregory (forthcoming) for a recent discussion and many citations.

causes higher consumption as time passes. In the Euler equation, it is the higher conditional variance of consumption that brings higher consumption growth. As the buffer stock of wealth rises, the conditional variance falls back and the extra consumption growth comes to an end.

Having isolated the effects of changing income uncertainty on consumption with greater specificity than in previous work, Carroll turns to the task of finding an observable measure of uncertainty. He chooses unemployment. His logic is that low-probability episodes of dramatic income loss are the most important kind of uncertainty and that unemployment is the reason for most such episodes. He makes the implicit assumption that periods of higher or rising unemployment are periods when the probability of large income losses from joblessness rises. The relation between aggregate unemployment and income volatility is more complicated. In a contraction, rates of job loss rise dramatically. However, when unemployment stops rising, unemployment typically remains elevated for several years. In fact, this appears to be the situation in the U.S. labor market since the middle of 1991. When unemployment is high but not rising, the rate of job loss is at roughly normal levels. The typical episode of joblessness lasts longer, however, because job-finding rates are lower in slack markets.

It is an interesting question whether a direct attack on the bufferstock hypothesis could be made using microeconomic panel data. Carroll comments skeptically on the interesting attempt by Karen Dynan to estimate an Euler equation with explicit treatment of the conditional variance of log consumption growth.² Her results lend little support to the buffer-stock hypothesis. I am not aware of any attempts to estimate structural consumption functions in microeconomic data with explicit uncertainty. It is possible that microeconomic panel data could be used to estimate the conditional distribution of future earnings and then to show that the dispersion of that distribution shifts the consumption function downward, as the buffer-stock hypothesis holds. In particular, panel data may make it possible to see directly how uncertainty changes from booms to recessions.

Carroll's applications of the model to consumption-saving puzzles are uniformly interesting, but not all are convincing. I found the discussion of the parallel behavior of consumption and income confusing. The

2. Dynan (1991).

joint behavior of consumption and income is the result of random shocks influencing both variables. The statement that the growth rate of income does not enter anywhere in the Euler equation misses the point. A random event that affects both income and consumption will have effects on all three terms in the Euler equation: the real interest rate, the conditional variance of future consumption, and the random increment. To decide if the parallel behavior of income and consumption is consistent with a particular theory, it is necessary to trace through the effects of an outside shock on the two variables to see whether they should move together for that type of shock. Carroll makes this kind of analysis for the buffer-stock regime, but he does not make a comparable analysis for the standard life-cycle regime. I do not think that he can reach a clear conclusion that parallel movements of consumption and income support the buffer-stock case over the standard case. Only a full general equilibrium analysis could resolve this question, but that analysis is way bevond the scope of this paper.

Carroll tackles the key question of the decline in consumption that occurs in recessions. His empirical results confirm that times when unemployment is high or rising are times when saving is higher than usual. Although Carroll is careful to consider the possibility that causality runs from higher saving to higher unemployment, rather than the other way around, I am not as persuaded as he is that the evidence favors causality from unemployment to saving. I find nothing to commend his identifying hypothesis that changes in preferences related to saving are uncorrelated with events in the economy six months earlier. To me, it is an open issue whether the story of the typical recession runs the way Carroll wants to tell it: some force depresses aggregate activity, raises unemployment, and thus raises income volatility, so consumption declines in relation to permanent income, thus worsening the recession. Although I see it as an open question whether Carroll's hypothesis is correct, I think the hypothesis deserves to be taken very seriously.

Carroll uses the buffer-stock hypothesis to explain at least part of the puzzle of low saving in the 1980s. Relaxation of borrowing constraints, improved credit reporting, and more widespread health and unemployment insurance are influences that depress saving in a buffer-stock world, but not in a life-cycle certainty-equivalence world. Carroll here offers an intellectually attractive alternative to the view that there was an unexplained spontaneous burst of extra consumption in the 1980s.

Carroll's paper is an important advance in consumption economics even after a decade of intense research. The paper also shows that much more remains to be done.

Stephen P. Zeldes: I will begin by describing how this paper fits into the existing literature on consumption and saving. Then I will briefly summarize what I think are the paper's main points. Finally, I will discuss some of the pitfalls and problems with Christopher Carroll's approach.

In recent years, there have been two main branches of the literature on consumption and saving. The first is what I will call the certainty or certainty equivalence (CEQ) approach; the second is the Euler equation approach. Let me describe each briefly.

Virtually all the work that has examined the optimal level of consumption and saving as a function of different variables has assumed either that labor income is variable, but nonstochastic, or has assumed quadratic preferences (so that consumption in the presence of uncertainty about labor income is identical to what it would be if labor income were set equal to its mean). The basic result is that optimal consumption is based on expected lifetime resources, calculated by adding financial wealth to the present value of expected future income. Consumption is set equal to the annuity value (possibly allowing for growth) of this summary statistic. The CEQ model is standard in almost all the literature on the life-cycle hypothesis (LCH) and permanent income hypothesis (PIH) that examines the level of consumption, and is the basis of virtually all of our intuition about consumption.

The second approach in the consumption literature has been the Euler equation approach, following the work of Robert Hall and many others.¹ This approach has focused on the expected growth rate of consumption, rather than on the level. The key result is that consumption growth should be orthogonal to lagged information. This approach typically assumes constant relative risk aversion (CRRA) preferences, although sometimes quadratic utility is assumed.² This approach has contributed substantially to our intuition about the expected growth rate of

2. Some recent work uses preferences that allow for a risk aversion parameter that is independent from the parameter for the intertemporal elasticity of substitution.

^{1.} Hall (1978).

consumption, but has not contributed to our intuition about the level of consumption.

Carroll's paper is part of a literature that combines parts of each of these approaches. This literature looks at the level of consumption and saving in the presence of random labor income and imperfect insurance, but uses preferences such as CRRA that are not immune to the effects of this income uncertainty. These preferences exhibit what Miles Kimball has termed "prudence": a positive third derivative of the utility function.³ These preferences generate a difference between what saving would be under certainty, with income set equal to its mean, and what saving is under uncertainty, a difference Havne E. Leland dubbed "precautionary saving."⁴ Leland and Kimball, as well as Agnar Sandmo, Truman Bewley, Jack Schechtman, Vera L.S. Escudero, and others, have been able to derive some properties of the solution analytically.⁵ Unfortunately, however, no one has been able to solve this optimal consumption problem analytically, so until recently economists were stuck. The advent of high-speed computers, however, has broken new ground on this problem. For example, in my 1989 paper, I applied some existing numerical dynamic programming techniques to the optimal consumption problem and examined the resulting numerical solution.⁶ I calculated the level of consumption as a function of wealth, the marginal propensity to consume out of wealth, and the expected growth of consumption over time as a function of initial wealth.

What has emerged from the papers mentioned above, and from work by Jonathan Skinner and Angus Deaton, as well as a number of papers in the consumption-based asset pricing literature, is a somewhat new intuition about optimal consumption under uncertainty.⁷ Although some of the intuition of the CEQ remains in these models, some surprising differences arise between consumption under uncertainty and under CEQ.

First, at any given wealth and income level, the level of consumption is lower (and thus saving is higher) under uncertainty than under certainty with income set to its mean. Second, at any given wealth and income level, the marginal propensity to consume out of wealth is higher

- 3. Kimball (1990a).
- 4. Leland (1968).
- 5. Sandmo (1970); Bewley (1977); Schechtman and Escudero (1977).
- 6. Zeldes (1989a).
- 7. Skinner (1988); Deaton (1991).

under uncertainty. While the result about the level of consumption depends on a positive third derivative of the utility function, the result about the marginal propensity to consume depends on derivatives higher than the third.

Third, the expected growth in consumption is higher under uncertainty. Fourth, all three of the above effects are stronger, the lower the level of financial wealth. Finally, the expected growth in consumption depends on the conditional variance of consumption growth.⁸

This brings us to Carroll's earlier work,⁹ some of which he reviews in this paper. One of the main results is that consumption more closely parallels income under uncertainty than under certainty.

Carroll examines both permanent and transitory differences in income growth. For comparisons of permanent differences over some range of growth rates, the model works similarly to a model with liquidity constraints. Given the rate of time preference, the interest rate, and the coefficient of relative risk aversion, if the income profile starts out steep enough, then steepening it further will cause the consumption profile to steepen; that is, wealth falls, causing expected consumption growth to rise. Carroll stresses this result in his other 1992 paper. I should point out that over lower ranges of the growth in income, this parallel breaks down; reducing the growth in income will cause the consumer to save more, but this will have only a small effect on the growth of consumption. Thus the consumption/income parallel occurs, but only over certain ranges of income growth.

Carroll also argues that a consumption/income parallel will occur for short-term fluctuations. This result builds on two existing ones. The first is the observation—made by many—that the conditional variance of consumption growth enters as a separate term in the linearized Euler equation. (This term is often assumed to be a constant in the empirical Euler equation consumption literature, although this is not the case in the consumption-based asset pricing literature.) The second is the result in my 1989 paper that the expected growth of consumption is significantly higher at low levels of wealth than at high levels of wealth. Carroll argues that low draws of income will lead to lower levels of wealth, and

9. Carroll (1992).

^{8.} Engen (1992) adds that the interest elasticity of saving can be much smaller in a precautionary saving model than in a certainty equivalence model.

that this leads to a higher conditional variance of consumption growth and higher mean consumption growth. This introduces the possibility that consumption can more closely track income in the short run, and that the commonly found violation of the linearized version of the Euler equation is due to this changing conditional variance. This is an intriguing theoretical possibility; the degree to which it is important in practice is still a wide-open question.

Deaton and then Carroll have also examined the implications for wealth accumulation over time.¹⁰ They have argued that in order to explain the relatively low wealth holdings of a large fraction of the population, it is necessary to assume that individuals have a relatively high rate of time preference or a high growth of income. In the model with CRRA preferences, a simple income process, and an infinite horizon, there is a single-state variable: the ratio of wealth to the permanent component of income. Because of the single-state variable property of the model, Deaton and Carroll have called the model with the higher rate of time preference a "buffer-stock model" of saving. Below some level of the single-state variable, the precautionary motive dominates and people accumulate wealth, on average. Above that level, the high rate of time preference dominates and people run down their wealth, on average. It is worth emphasizing that this buffer-stock model is a standard model of intertemporal consumer optimization under uncertainty, with a parameterization such that consumers would want to borrow in the absence of uncertainty.

After deriving analytic and numerical results, the next step is to see whether precautionary saving is important in practice, and whether a model with precautionary saving can explain empirical phenomena that are puzzling from the point of view of a CEQ model. There are two ways of doing this: the first uses numerical methods, simulation, and/or calibration to see whether an appropriately parameterized model can match certain stylized facts about the world. The second approach is to perform empirical estimation and test directly for the presence of precautionary effects. A number of very recent papers do this, including Karen Dynan's.¹¹ The current Carroll paper uses both these approaches; the numerical work guides the empirical work. It examines the effects of

10. Deaton (1991); Carroll (1992).

11. Dynan (1991).

changes in expected income growth, as well as changes in uncertainty about future income. A main result is that small changes in the probability of extremely low realizations of income can have significant effects on precautionary saving. Carroll interprets these low realizations of income as spells of unemployment.

I have only one comment on the specifics of the numerical results. The comparative statics exercises done in the paper are suggestive, but do not completely correspond to the types of changes households face in the world. For example, consider the effects of a change in the growth rate of income or a permanent change in uncertainty. The paper assumes that individuals have always faced-and always will face-that same growth rate or that same uncertainty. The paper then changes the growth rate or uncertainty and examines the dynamic effects. That is, people face a change to which they had ex-ante assigned zero probability. While this approach used to be a commonly performed comparative statics exercise, it is an inappropriate experiment to perform. It is especially problematic here, where the whole point of the exercise is that people are forward-looking and take into account all possible contingent outcomes. The correct way to do the calculations (which, granted, is very difficult) is to let people face a probability distribution for future income growth and possibly for uncertainty, and then examine the optimal response to various realizations. This is an important area for future research.

Let me now turn to the empirical work. The goal of Carroll's paper is to examine the empirical effects of an economy-wide change in the degree of uncertainty facing individuals. Much of what is new in the paper is the direct empirical work, and in particular the idea of using the unemployment rate (or a proxy of it) as a measure of this uncertainty. Carroll claims that a high (or rising) unemployment rate corresponds to a high level of income uncertainty. His empirical work examines the relationship between the saving rate or consumption growth and a proxy for unemployment expectations.

Let me briefly describe the variables used in the empirical work. The paper uses answers to survey questions about saving, changes in aggregate unemployment, and changes in individual income. It is worth reviewing exactly what these questions are.

The saving question asks: "If there were a major purchase that you wanted to make, do you think that now is a time when it would be okay to

use some of your savings, or is now a time when you would be especially reluctant to use some of your savings?" Carroll constructs a measure (SAV) equal to the fraction of the people who answered that they would be reluctant minus the fraction who answered that it would be okay.

I have two comments about this variable. First, because it refers to a "major purchase," it is clearly asking about durables. However, all the theory in the paper is about nondurables and services. There are undoubtedly important links between uncertainty, illiquidity, and durables expenditures (some of which have been modeled in other recent papers),¹² but they are not modeled or carefully discussed in this paper. This makes me question the validity of the mapping from the theory to the empirical work.

Second, I can not tell from this question what the suggested alternative is. If respondents answer that they would be reluctant to use some of their savings now, does that mean that they would be unlikely to make the major purchase or that will they find the money elsewhere (by using current income or cutting other expenditures) and purchase it anyway? It is not clear what this question reveals.

Another question is about unemployment changes. It asks: "How about people out of work during the next 12 months—do you think that there will be more unemployment than now, about the same, or less?" Again, Carroll uses a constructed variable (MU): the fraction of households who believe unemployment will increase minus the fraction who believe it will decrease. I also have two comments about this variable. First, it is about the change in unemployment, not the level. Second, it measures how much of a consensus exists about the change in unemployment, not the amount that unemployment is expected to rise. Thus if a large fraction of the population anticipates a very small rise in unemployment, MU will be very high, even though the expected change in unemployment may be very low. Calling the variable the "expected change in unemployment" is therefore misleading.

Yet another question has to do with the change in income. It asks: "[In] the next year or two, do you expect that your (family) income will go up more than prices will go up, or less than prices will go up?" Carroll uses the constructed variable (MY): the fraction of households who believe their real household income will increase minus the fraction who

12. See, for example, Eberly (1992).

146

believe it will decrease. This variable has the same problem as MU: it does not capture the degree to which income is expected to rise, but rather measures the fraction of people who expect their income to rise. Also note—and this is what is intended—that this question is asking about individual income, while the previous question was asking about aggregate unemployment.

I will focus on four sets of tests and results: those of tables 5, 6, 7, and 9. I will first summarize the intuition behind the tests and the results in all four tables, and then comment on them.

The first part of table 5 regresses the first variable described, SAV, on the two expectations variables, MU and MY. The idea is as follows: if MU is high, lots of people expect unemployment to rise, so that uncertainty about who will become unemployed is high, which means that precautionary saving is high; therefore now is a bad time to buy. This suggests a positive coefficient on MU, and Carroll does find a significant positive coefficient.

Table 6 regresses the personal saving rate on MU, MY, and the unemployment rate, using quarterly data and (in some cases) instrumental variables. The idea is similar to the test in table 5, except here the actual saving rate, s, is used rather than the survey variable SAV, and the level of unemployment is included as well. An expected increase in unemployment should correspond to high uncertainty and, thus, to a high saving rate. The empirical results are mixed, but, in some cases, MU enters with the correct sign, and in others the unemployment rate itself enters with a positive sign.

Table 7 reports a regression like the ones run by John Campbell and N. Gregory Mankiw.¹³ The change in consumption is regressed on the (instrumented) actual change in income and MU. The test is as follows: once MU is included as a regressor, is there any remaining evidence that some households follow rule-of-thumb behavior of consuming their income in every period? If so, then after allowing for some rule-of-thumb consumers, does MU help predict consumption growth? The standard theory of precautionary saving implies that MU should enter with a positive sign: higher MU implies higher uncertainty, which implies a faster expected growth of consumption. Carroll finds that including MU decreases the sign of the coefficient on the growth in income, but it remains

^{13.} Campbell and Mankiw (1989).

significantly greater than zero. He also finds that *MU does* seem to have a separate role, and enters the equation significantly; unfortunately, the sign of the coefficient is negative rather than positive. So, Carroll says that in order to reconcile this finding with the theory, the theory must be augmented in one of two ways: by adding habit formation or adjustment costs for consumption.

Finally, in table 9, Carroll takes up the current recovery. He examines the consumption growth forecast errors from the table 7 equations (the Campbell/Mankiw model with and without MU) for the period since the 1990:3 business cycle peak, and also for the period since his estimate (1991:2) of the business cycle trough. He finds that there is slightly less underprediction of the growth of consumption when MU is included than when it is not. His interpretation is that high uncertainty has generated low consumption growth in the current recovery.

Let me comment on these tests and results. Overall, I find the empirical results intriguing and suggestive, but they do not constitute strong evidence supporting the theory. A large gap between theoretical results and the empirical tests often exists that requires a leap of faith on the part of the reader. Carroll provides only a sketch of how the empirical tests might follow from the theory; much more work is needed to provide the bridge and the road map between the two.

The other main difficulty that I have with the empirical implementation is that I am not convinced that the survey variable MU is a good proxy for future income uncertainty. Two links must be made to justify MU as a good proxy. One is the link from MU to expected changes in unemployment. The other is the link from expected changes in unemployment to uncertainty. I have already discussed my reservations about the MU variable as a proxy for the expected change in unemployment. As far as the second link, Carroll does not make it clear whether income uncertainty is high when the expected change in unemployment is high or when the expected *level* of unemployment is high: that is, is uncertainty higher if unemployment is expected to rise from 6 to 8 percent next period or if unemployment is expected to stay constant at 10 percent? Much of the discussion in the paper is about the level of unemployment, but MU seems to be serving as a proxy for the expected change in unemployment. Presumably, whether it is expected levels or changes (or neither) that is relevant for uncertainty has to do with the

148

dynamics of unemployment at the household level: how does an individual's probability of job loss change with the aggregate unemployment rate? Unfortunately, the paper does not include a discussion of, or evidence about, labor market dynamics. Direct evidence on the relationship between the aggregate unemployment rate and the probability and duration of individual spells of unemployment should be a key goal of future research in this area.

Finally, the MU variable could easily be a proxy, in my mind, for the mean of expected future income, rather than the variance. Carroll tries to adjust for this by including the MY variable in the regressions, arguing that MY captures the expected change in income, and any additional information in MU must be about other moments of the distribution. However, because MY is the fraction of the population expecting increases in income, rather than the actual expected income itself, it seems unlikely that it is fully capturing the expected change in income. If a large fraction of the population believes unemployment will be rising, this corresponds to an expected drop in the mean of future income. The standard certainty equivalence model (with no precautionary saving) would predict a high saving rate. Therefore I am not convinced that the tests are capturing a precautionary saving effect, rather than an effect that would arise in a standard certainty equivalence permanent income/life-cycle model.

Thus while the evidence linking saving to MU is suggestive that precautionary saving may be important, more work (both theoretical and directly empirical) is needed to show that MU is indeed capturing household uncertainty about future income. This paper is likely to stimulate further research in this area.

In conclusion, I think that this paper contributes to the growing literature on the effects of uncertainty on optimal saving and consumption, and in particular draws out many of the macroeconomic implications of the theory. Some of the numerical exercises help change our intuition about saving under uncertainty. The most innovative part of the paper is the idea of using MU as an empirical proxy for individuallevel uncertainty about income. I find the empirical results intriguing and, on the whole, intuitively appealing. However, I am not fully convinced that they actually provide support for this precautionary saving model.

General Discussion

Olivier Blanchard noted that the common empirical finding of a low interest elasticity of saving is consistent with the paper's result that the target wealth ratio is unlikely to be sensitive to changes in the interest rate. However, Alan Blinder noted that interest rates could have important wealth effects that would affect saving in the model. If interest rates rise, wealth is reduced; this, according to the model, raises the conditional variance of consumption and consumption growth. In response, consumers would lower current consumption and increase saving. In this vein, Richard Cooper asked whether the increase in wealth in the 1980s could be a reason for the lower saving rate of that period. Carroll noted that a large fraction of personal wealth is held by the very wealthy, people to whom he does not claim his model applies. He further argued that macroeconomic models of consumption which explicitly incorporate an effect of wealth on the saving rate are unable to explain the magnitude of the decline in the saving rate, so the wealth effect cannot be the only explanation for the decline in saving.

Christopher Sims observed that using models of this type to explain the business cycle is risky unless they are embedded in a general equilibrium framework. Interest rates and income growth are not exogenous as assumed in the model. Sims noted, for example, that neoclassical growth models can alter standard consumption growth results which assume an exogenous constant interest rate.

Blinder observed that the model also suggests that a widening disparity in incomes may affect saving. Increasing inequality of income and wealth is likely to involve an increase in the probability that some consumers might have zero incomes. Because buffer-stock saving is especially sensitive to the risk of very low incomes, the model predicts that growing inequality would lead to lower aggregate consumption levels.

Blanchard noted that the model as specified does not capture the lifecycle motivation of saving for retirement; income is expected to grow over the lifetime and, indeed, is assumed to continue growing until death. In real life, consumers must increase their target wealth ratio before they retire to provide for retirement. He suggested interpreting the model as part of a two-track saving process in which individuals make contractual commitments for their retirement. Once consumers have

made such commitments to cover major saving for retirement, then they may behave according to the buffer-stock model with the remainder of their incomes. He suggested a possible rationale for this sort of twotrack saving behavior: nonexponential discounting by consumers that reflected high levels of impatience in the short run, but less impatience over the longer run. Carroll acknowledged that the version of his model with constant income growth until death does not provide a role for lifecycle saving. He added that if the model is modified so that incomes grow rapidly during the early part of the life cycle and then turn down, his model would predict buffer-stock saving early in the life cycle; but that later in life, as life-cycle considerations dominate, buffer-stock considerations would fade.

Gregory Mankiw was skeptical that the precautionary saving effects predicted by the model could be seen in the data. He noted that recent research by Dynan (1991) using panel data finds little evidence of a correlation between consumption growth and the variance of consumption. But he also suggested how this finding might not be inconsistent with the model: people may self-select into occupations that have different income risks, on the basis of their degree of risk aversion. This self-selection process would obscure the evidence of the effects predicted by the model in panel data.

Julio Rotemberg cautioned against interpreting the estimates from aggregate Euler equations in terms of the characteristics governing individual behavior in the model. As aggregate unemployment varied cyclically, the probability of becoming unemployed would vary in different ways for different individuals. For many, the risk of job loss would not change perceptibly while for others it would change dramatically.

Robert Gordon distinguished between two descriptions of the current recession: the popular account of a "hangover recession" and the account proposed by the model. On the first description, one of the current "hangovers" stems from an excessive increase in debt by households and firms during the 1980s, which must now be serviced and retired. By contrast, the Carroll model could imply that no hangover exists because greater indebtedness resulted from an equilibrium adjustment by consumers to loosened borrowing constraints. Carroll responded that the "hangover" description of the current recession could still be accurate because an increase in uncertainty could raise consumers' desired target wealth, making them uncomfortable holding debts that were viewed as

tolerable when times were better. He pointed out that, in this interpretation, the increase in unemployment expectations is what converted an acceptable debt burden in 1989 into a debt "hangover" in 1991. Sims cited reductions in the penalties for personal bankruptcy during the 1980s as another equilibrium explanation for consumers' greater willingness to borrow.

Martin Baily related the model to inventory theories of the firm. If firms, like Carroll's consumers, are concerned about low-probability but disastrous events such as bankruptcy, then the use of inventories for production smoothing may not occur over the business cycle. Rather, firms would run down inventories as the economy enters recession—as is in fact observed. However, Baily cautioned that people often do not behave rationally in planning for low-probability disastrous events, citing the apparently irrational behavior in the market for flood or earthquake insurance. Such irrational behavior would weaken the link between current consumption and the probability of future unemployment. Carroll noted the ample research into this type of behavior. His impression of the literature was that people planned incorrectly, not that they failed to plan altogether. Thus, he argued, a link between unemployment expectations and consumption may exist, even if those expectations do not reflect an optimal forecast of unemployment.

Drawing an analogy to physics, William Nordhaus noted that, like the physical anomalies that appear at low temperatures, strange behavior which the model does not capture occurs at the low end of the income scale.He cited several examples of consumers whose behavior is not easily explained by the buffer-stock model: students, welfare recipients, and business people who go on a spending spree as personal bankruptcy looms. Baily argued for the importance of contagion effects: consumers, in their spending patterns, may be influenced by the attitudes and behavior of their reference group. When the prevailing consumer mood is generally depressed during a recession, consumers may spendless, regardless of the true expectation of their own future incomes. Carroll acknowledged the possibility of contagion effects, and suggested that beliefs about the increased risk of unemployment may be important for explaining consumption even if economists have trouble detecting the marginal effect of increased unemployment on the conditional distribution of income. He argued that consumers should act on the basis of their own beliefs, whether or not those beliefs can be shown to be optimal forecasts.

References

- Abowd, John M., and David Card. 1989. "On the Covariance Structure of Earnings and Hours Changes." *Econometrica* 57 (2):411–45.
- Auerbach, Alan J., and others. 1992. "The Increasing Annuitization of the Elderly." Ruth Pollack Working Papers Series on Economics, No. 6. Boston University.
- Backus, David K., and Allan W. Gregory. Forthcoming. "Theoretical Relations between Risk Premiums and Conditional Variances." *Journal of Business* and Economic Statistics.
- Bernheim, B. Douglas, and John B. Shoven. 1988. "Pension Funding and Saving." In *Pensions in the U.S. Economy*, edited by Zvi Bodie, John B. Shoven, and David A. Wise. Chicago: University of Chicago Press.
- Bewley, Truman. 1977. "The Permanent-Income Hypothesis: A Theoretical Formulation." Journal of Economic Theory 16(2):252–92.
- Blinder, Alan S., and Angus S. Deaton. 1985. "The Time Series Consumption Function Revisited." *BPEA*, 2:1985, 465–511.
- Bound, John, and Alan B. Krueger. 1989. "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?" Unpublished paper. Cambridge, Mass.: National Bureau of Economic Research (August).
- Bovenberg, A. Lans, and Owen Evans. 1989. "National and Personal Saving in the United States: Measurement and Analysis of Recent Trends," Working Paper 103. Washington, D.C.: International Monetary Fund (December).
- Burtless, Gary. "Why Is Insured Unemployment So Low?" BPEA, 1:1983, 225-49.
- Caballero, Ricardo J. 1991. "Earnings Uncertainty and Aggregate Wealth Accumulation." *American Economic Review* 81:859–71.
 - ——. 1992. "Near-Rationality, Heterogeneity, and Aggregate Consumption," Working Paper 4035. Cambridge, Mass.: National Bureau of Economic Research (March).
- Campbell, John Y. 1987. "Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis." *Econometrica* 55(6):1249–73.
 - ——, and Angus S. Deaton. 1989. "Why Is Consumption So Smooth?" *Review of Economic Studies* 56:357–74.
 - ——, and N. Gregory Mankiw. 1989. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." In *NBER Macroeconomics Annual 1989*, edited by Olivier Blanchard and Stanley Fischer. Cambridge, Mass.: MIT Press.
- Canner, Glenn B., and Charles A. Luckett. 1992. "Developments in the Pricing of Credit Card Services." Federal Reserve Bulletin 78 (September): 652–66.
- Carroll, Christopher D. 1992. "Buffer Stock Saving and the Permanent Income Hypothesis." Unpublished paper. Washington, D.C.: Federal Reserve Board.

- —, and Lawrence H. Summers. 1991. "Consumption Growth Parallels Income Growth: Some New Evidence." In *National Saving and Economic Performance*, edited by B. Douglas Bernheim and John B. Shoven. Chicago: University of Chicago Press.
- ——, and Andrew Samwick. 1992. "The Nature and Magnitude of Precautionary Wealth." Unpublished paper. Washington, D.C.: Federal Reserve Board (March).
- Cochrane, John H. 1991. "Univariate vs. Multivariate Forecasts of GNP Growth and Stock Returns: Evidence and Implications for the Persistence of Shocks, Detrending Methods, and Tests of the Permanent Income Hypothesis." Unpublished paper. University of Chicago (February).
- Constantinides, George M. 1990. "Habit Formation: A Resolution of the Equity Premium Puzzle." *Journal of Political Economy* 98(3):519–43.
- Cuthbertson, Keith, and David Barlow. 1991. "Disequilibrium, Buffer Stocks, and Consumers' Expenditure on Non-Durables." *Review of Economics and Statistics* 1991(4):643–53.
- Deaton, Angus S. 1991. "Saving and Liquidity Constraints." *Econometrica* 59(5):1221-48.
- Duesenberry, James S. 1949. Income, Saving, and the Theory of Consumer Behavior. Cambridge, Mass.: Harvard University Press.
- Duncan, Greg J., and Daniel H. Hill. 1985. "An Investigation of the Extent and Consequences of Measurement Error in Labor-economic Survey Data." *Journal of Labor Economics* 3(4):508–32.
- Dynan, Karen E. 1991. "How Prudent Are Consumers?" Unpublished paper. Harvard University (December).
- Eberly, Janice C. 1992. "Flexible (*S*,*s*) Bands, Uncertainty, and Aggregate Consumer Durables." Unpublished paper. Rodney L. White Center for Financial Research, the Wharton School, University of Pennsylvania (July).
- Engen, Eric. 1992. "Precautionary Saving, Consumption, and Taxation in a Life-Cycle Model with Stochastic Earnings and Mortality Risk." Ph.D. diss. University of Virginia (January).
- Epstein, Larry G., and Zin, Stanley E. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns I: A Theoretical Framework." *Econometrica* 57(4):937–69.
- Feldstein, Martin. 1974. "Social Security, Induced Retirement, and Aggregate Capital Accumulation." Journal of Political Economy 82(5):905-26.
- Flavin, Marjorie A. 1981. "The Adjustment of Consumption to Changing Expectations About Future Income." *Journal of Political Economy* 89(5):974–1009.
- Friedman, Milton. 1957. A Theory of the Consumption Function. Princeton, N.J.: Princeton University Press.
- Fuhrer, Jeffrey C. 1992. "Do Consumers Behave as the Life-Cycle/Permanent-Income Theory of Consumption Predicts?" *New England Economic Review* (Sept./Oct.):3–14.
- Hall, Robert E. 1978. "Stochastic Implications of the Life Cycle-Permanent In-

come Hypothesis: Theory and Evidence." Journal of Political Economy 86(6):971-87.

—. 1988. "Intertemporal Substitution in Consumption." *Journal of Political Economy* 96(2):339–57.

—, and Frederick S. Mishkin. 1982. "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." *Econometrica* 50(2):461–81.

- Health Insurance Association of America. 1987. 1986–1987 Source Book of Health Insurance Data. Washington, D.C.: Health Insurance Association of America.
- Heaton, John. 1990. "The Interaction Between Time-nonseparable Preferences and Time Aggregation," Working Paper 3181-90-EFA. Cambridge, Mass.: Massachusetts Institute of Technology (June).
- Kimball, Miles S. 1990a. "Precautionary Saving in the Small and in the Large." *Econometrica* 58(1):53–73.

——. 1990b. "Precautionary Saving and the Marginal Propensity to Consume." Working Paper 3403. Cambridge, Mass.: National Bureau of Economic Research (July).

- Leland, Hayne E. 1968. "Saving and Uncertainty: The Precautionary Demand for Saving." *Quarterly Journal of Economics* 82(3):465-73.
- MaCurdy, Thomas E. 1982. "The Use of Time Series Processes to Model the Error Structure of Earnings in Longitudinal Data Analysis." *Journal of Econometrics* 18(1):83–114.
- Mankiw, N. Gregory. 1981. "Hall's Consumption Hypothesis and Durable Goods." Journal of Monetary Economics 10:417–25.

------, and Stephen P. Zeldes. 1990. "The Consumption of Stockholders and Non-Stockholders." Working Paper 3402. Cambridge, Mass.: National Bureau of Economic Research (July).

- Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55(3):703–08.
- Pagano, Marco, and Tullio Jappelli. 1991. "Information Sharing in Credit Markets." Discussion Paper 579. London: Centre for Economic Policy Research (October).
- Sandmo, Agnar. 1970. "The Effect of Uncertainty on Saving Decisions." *Review of Economic Studies* 37(3):353-60.
- Schechtman, Jack. 1976. "An Income Fluctuation Problem." Journal of Economic Theory 12:218–41.

-----, and Vera L.S. Escudero. 1977. "Some Results on 'An Income Fluctuations Problem." *Journal of Economic Theory* 16(2):151-66.

- Skinner, Jonathan. 1988. "Risky Income, Life-cycle Consumption, and Precautionary Savings." Journal of Monetary Economics 22(2):237-55.
- Stiglitz, Joseph, and Andrew Weiss. 1981. "Credit Rationing in Markets with Imperfect Information." American Economic Review 71:393-410.
- Stock, J. H., and M. W. Watson. 1989. "New Indexes of Leading and Coinci-

dent Economic Indicators." In NBER Macroeconomic Annual 1989, edited by Olivier Blanchard and Stanley Fischer. Cambridge, Mass.: MIT Press.

- Summers, Lawrence H. 1981. "Capital Taxation and Accumulation in a Life Cycle Growth Model." *American Economic Review* 74(4):533–45.
 - ——, and Christopher D. Carroll. "Why Is U.S. National Saving So Low?" BPEA, 2:1987, 607–35.
- Thaler, Richard H, and H. M. Shefrin. 1981. "An Economic Theory of Self-Control." Journal of Political Economy 89(2):392–406.
- Touche Ross. 1989. The Touche Ross Personal Financial Management and Investment Workbook. Englewood Cliffs, N.J.: Prentice-Hall.
- University of Michigan. Survey Research Center. Institute for Social Research. Wave XVIII. A Panel Study of Income Dynamics. Ann Arbor, Mich.
- University of Michigan. Survey Research Center. Institute for Social Research. 1992. Surveys of Consumers. Ann Arbor, Mich. (June).
- U.S. Department of Labor. Various years. Unemployment Insurance: State Laws and Experience (UI Chartbook).
- Wilcox, David W. 1991. "Household Spending and Saving: Measurement, Trends, and Analysis." *Federal Reserve Bulletin* 77(January):1-17.
- Zeldes, Stephen P. 1989a. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *Quarterly Journal of Economics* 104:275–98.
 - ———. 1989b. "Consumption and Liquidity Constraints: An Empirical Investigation." Journal of Political Economy 97:305–46.