This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: NBER Macroeconomics Annual 1997, Volume 12

Volume Author/Editor: Ben S. Bernanke and Julio Rotemberg

Volume Publisher: MIT Press

Volume ISBN: 0-262-02435-7

Volume URL: http://www.nber.org/books/bern97-1

Publication Date: January 1997

Chapter Title: Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets

Chapter Author: Chris Carroll, Wendy Dunn

Chapter URL: http://www.nber.org/chapters/c11039

Chapter pages in book: (p. 165 - 230)

Christopher D. Carroll and Wendy E. Dunn
THE JOHNS HOPKINS UNIVERSITY

Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets

1. *Introduction*

The U.S. recession that began in 1990 and the feeble recovery that followed differed from the pattern of previous postwar business cycles in several respects, most notably in the sustained weakness in consumption spending, particularly for durable goods. Blanchard (1993) estimates a simple macroeconomic model and finds that the recession was largely the result of a "consumption shock." Hall (1993) finds an important role for a "spontaneous decline in consumption," especially for durable goods. Furthermore, structural macroeconomic models like the FRB-US model substantially overpredicted consumption spending throughout the 1990 recession and especially the early recovery period.

In December 1991, as the economy struggled to make its way out of recession, Federal Reserve Chairman Alan Greenspan included the following statements in Congressional testimony on the state of the economy:

During the 1980s, large stocks of physical assets were amassed in a large number of sectors, largely financed by huge increases in indebtedness. . . . In the house-

We are grateful to Carl Christ, Darrel Cohen, Angus Deaton, Karen Dynan, Bruce Fallick, James Fergus, Michael Fratantoni, Simon Gilchrist, Spencer Krane, Sydney Ludvigson, Greg Mankiw, Randy Mariger, Louis Maccini, David Reifschneider, Martha Starr-McCluer, Daniel Sichel, and seminar participants at Johns Hopkins University, the Board of Governors of the Federal Reserve System, and the NBER Macroeconomics Annual Conference for valuable comments and all sorts of help. The set of RATS programs that generate all of our empirical results, the set of Mathematica programs that generate all of our theoretical results, and a companion paper describing details of our empirical and theoretical methodology are all available at Carroll's home page, http://www.econ.jhu.edu/ccarroll.

hold sector, purchases of motor vehicles and other consumer durables ran for several years at remarkably high levels and were often paid for with installment or other debt that carried extended maturities. In some parts of the United States, the household spending boom reached to the purchase of homes. . . . The aftermath of all this activity is a considerable degree of financial stress in the household sector. (Greenspan, 1992)

In this testimony and elsewhere, Greenspan consistently blamed the 1990–1991 recession and the subsequent painfully slow recovery on the "deteriorated balance sheets" of both firms and households resulting from the buildup of debt in the 1980s. Figure 1 shows that the runup in household debt in the 1980s was indeed impressive. Most of this growth was in mortgage debt, spurred by the financial deregulation of the early 1980s which led to low down-payment requirements on home purchases.

The problematic part of what we will call the *Greenspan hypothesis* is that it provides no explanation for why balance-sheet positions that consumers voluntarily chose in the spring and summer of 1990 were

Figure 1 DEBT-TO-INCOME RATIO

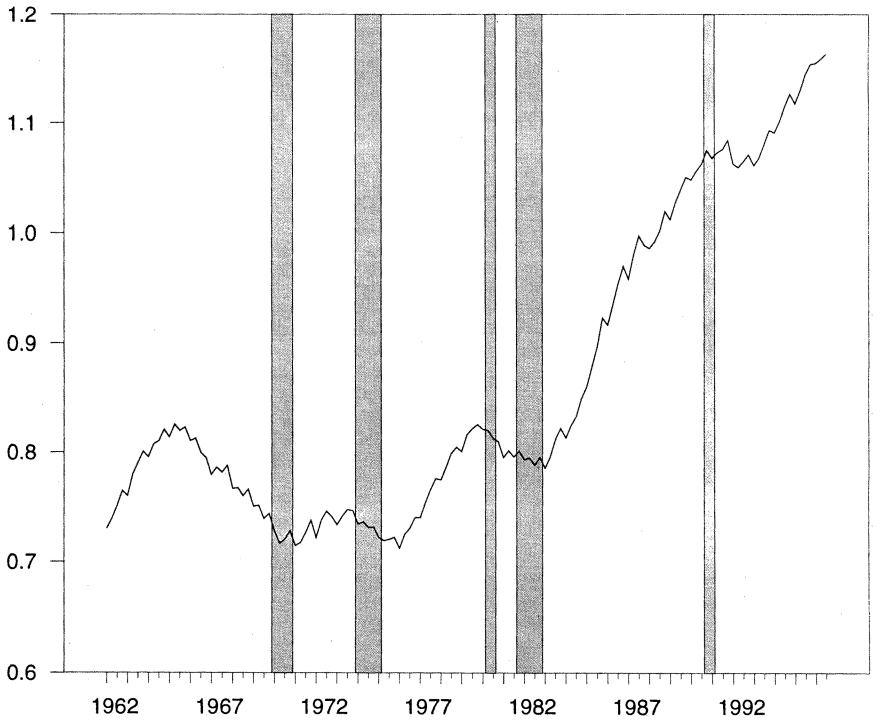

Figure 2 UNEMPLOYMENT EXPECTATIONS

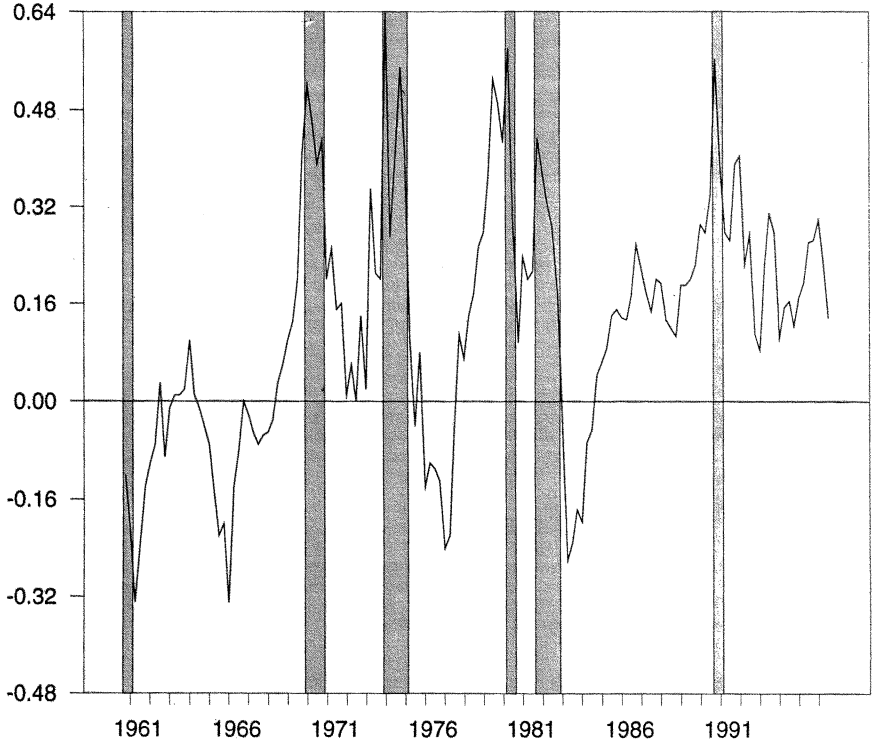

suddenly a major contractionary force in the fall of 1990 and in 1991. One plausible possibility is that an aggregate "target" consumer balance-sheet position depends, among other things, on the degree of consumers' uncertainty about the future, and in particular on their perceptions about the risk of future unemployment spells. Figure 2 plots the best available data on household unemployment expectations, from the University of Michigan's monthly surveys of consumers.[1] Unemployment expectations deteriorated sharply in the fall of 1990, right at the time of the "spontaneous" consumption drop.[2] The natural interpretation is that

1. The index is equal to the fraction of consumers surveyed who thought unemployment would rise over the next twelve months minus the fraction who thought unemployment would fall.
2. We choose this unemployment expectations index to measure consumer sentiment for several reasons. First, it has a much clearer definition than the more commonly used overall measures of sentiment, which combine in arbitrary ways the answers to questions about the past, present, and future conditions in a variety of largely unrelated markets. Second, one of the principal theoretical results in the precautionary-saving literature is that large shocks like unemployment spells should be disproportionately

it was the deterioration in unemployment expectations that converted a balance-sheet position which consumers had voluntarily chosen in happier times into one that required serious "repair." Indeed, it might appear tempting to attribute the consumption drop in 1990 entirely to the deterioration in sentiment and to dismiss the condition of household balance sheets as a sideshow.[3] One difficulty of this interpretation, however, is that unemployment expectations always deteriorate near the beginning of a recession (see Figure 2 again), and the 1990 experience was not sufficiently different from previous recessions to explain why consumption growth was weaker than it usually is during recessions. The behavior of the unemployment expectations index was more unusual after the trough of the recession; usually the index plummets just after the trough, but unemployment expectations remained quite high for a long time after the 1991 trough.[4] Still, even consumption models which incorporate the unemployment expectations index have large negative residuals during and after the 1990 recession, implying that the consumption weakness cannot be explained as simply reflecting consumer pessimism.

Prompted by this debate, this paper is a broad attempt to make sense of the relationship between household balance sheets, unemployment expectations, and household purchases. We begin (in Section 2) by documenting what we take to be the main stylized facts about the empirical relationships between consumer purchases, household balance sheets, and uncertainty. The only systematic relationship we are able to uncover between balance-sheet measures and spending is a robust *positive* correlation between lagged debt growth and the current level of spending on durables, a relationship which is most easily interpreted as reflecting simultaneity rather than a causal link. However, we do identify another empirical regularity: our preferred measure of uncertainty, the lagged value of the unemployment expectations index plotted in Figure 2, is robustly correlated with every measure of consumer spending, even after controlling for permanent income as best we can (and in particular after controlling for whatever information unemployment expectations contain about future income).

---

important in determining behavior as compared with small shocks such as wage fluctuations for employed consumers. Finally, the unemployment expectations index is considerably more robustly correlated with most measures of spending than are overall sentiment measures.

3. Both Blanchard (1993) and Hall (1993) suggest that the decline in sentiment was important, but neither emphasizes balance-sheet issues.
4. It is interesting to note that the index was right, in the sense that the unemployment rate did remain unusually high for an unusually long period after the trough.

With these results in mind, we then (in Section 3) construct a theoretical model of the durable-goods purchase problem for consumers who face the possibility of unemployment spells. Because analytical solutions are not available when there is labor-income uncertainty, we solve the model numerically. We find that the model implies that a rise in uncertainty causes consumers to delay durables purchases [formally, the lower trigger of the (S,s) rule jumps down; hence our title]. We then compare simulation results from the model with our empirical evidence for the U.S. economy, and find that the model explains some but not all of the empirical findings. In particular, the model implies a much stronger role for changes in unemployment expectations, and a weaker role for the lagged level of unemployment expectations, than we find in the data. Finally, in Section 6, we show that the model implies that a financial liberalization which loosens liquidity constraints will cause a runup in aggregate debt like the runup shown in Figure 1, and that in the liberalized economy the reaction of durables purchases to uncertainty is intensified. Thus our model potentially rationalizes the idea that the runup of consumer debt in the 1980s was partly responsible for the puzzling weakness of consumption spending during and after the 1990 recession. Furthermore, the model implies that the continuing growth of the debt ratio may be making consumption increasingly vulnerable to swings in consumer sentiment.

2. *Empirical Results*

2.1 BALANCE SHEETS AND NONDURABLES CONSUMPTION GROWTH

Although housing and other durable goods account for most of the volatility of consumption spending over the business cycle, we begin our empirical work by examining spending on nondurable goods. Partly this is because virtually no existing work has examined the effect of either balance sheets or time-varying unemployment expectations on nondurables spending, and these are important questions in their own right. Partly, we examine nondurables because one of the innovations of our theoretical model is our joint treatment of durables and nondurable goods.[5] Thus, in principle, even in the absence of time-varying unemployment risk our theoretical model might generate different predictions for nondurables spending than standard models.

The benchmark model with which we intend to compare both empirical

5. Most previous modeling efforts, with the exception of Bernanke (1985), have assumed utility flows either solely from nondurables or solely from durables, or at the very least that utility from durables and nondurables is separable.

results and the theoretical predictions of our model is the representative-agent, certainty-equivalent version of the permanent-income model (henceforth, CEQ PIH model), as used, for example, by Campbell (1987), Campbell and Deaton (1989), and many others. In this model, consumption is equal to *permanent income*, defined as the annuity value of total wealth, human and nonhuman:

$$C_t = \frac{r}{1+r}(W_t^h + W_t^n),$$

$$W_t^h = \sum_{s=t}^{\infty} \left(\frac{1}{1+r}\right)^{s-t} Y_s,$$

where Y_s is total noncapital income (labor income plus net transfers) in period s. We define a variable which we will call *annuity labor income* A_t as the annuity value of human wealth[6]:

$$A_t = \frac{r}{1+r} W_t^h.$$

As Hall (1978) famously pointed out, one of the implications of this model is that lagged information should have no predictive power for current consumption growth. Campbell and Mankiw (1989) showed that all of the failures of the aggregate CEQ PIH model could be explained by a model in which a fraction λ of aggregate labor income goes to rule-of-thumb consumers who simply spend all available income in each quarter while $1 - \lambda$ of income accrues to consumers who behave according to the CEQ PIH model. These assumptions, plus a few approximations, lead to an estimating equation of the form

$$\Delta \log C_t = \gamma_0 + \gamma_1 E_{t-1} \Delta \log Y_t + \epsilon_t,$$

where the expectation is taken with respect to a set of instruments dated $t - 1$.[7] Because, strictly speaking, the model applies only to the consump-

6. We adopt this terminology partly to avoid confusion between the variable in this model and the permanent-labor-income variable in our theoretical model.
7. Because time aggregation can introduce an MA(1) error term, the usual procedure is to use instruments dated $t - 2$. However, as Carroll, Fuhrer, and Wilcox (1994) argue, this unnecessarily discards potentially valuable information in variables dated $t - 1$. We follow those authors in pursuing a nonlinear estimation methodology that allows us to use instruments dated $t - 1$ and to impose the orthogonality restriction directly. Our instruments for income growth are the same as those used by Carroll, Fuhrer, and Wilcox (1994): three lags each of income growth, consumption growth, the change in the

Table 1 THE SENTIMENT-AUGMENTED CAMPBELL–MANKIW MODEL[a]

Row	*Balance-sheet measure*	$E_{t-1}\Delta \log Y_t$	UE_{t-1}	*Balance sheet*	θ	SSR	D–W
1		0.509 (4.13)***			0.086 (0.93)	0.49	1.98
2			−1.310 −(3.69)***		0.136 (1.47)	0.58	1.97
3		0.269 (1.64)	−0.906 −(2.18)**		0.092 (0.99)	0.50	1.98
4	$\Delta \log D_{t-1}$	0.246 (1.50)	−0.690 −(1.55)	0.095 (1.33)	0.088 (0.94)	0.49	2.00
5	rD_{t-1}/Y_{t-1}	0.257 (1.57)	−0.820 −(1.90)*	−0.073 −(0.93)	0.0937 (1.00)	0.49	1.98
6	D_{t-1}/A_{t-1}	0.247 (1.45)	−0.906 −(2.15)**	−0.002 −(0.33)	0.096 (1.02)	0.50	1.97

[a]Dependent variable is nondurable consumption growth, quarterly data, 1963:3–1994:3. *, significant at 10% or better; **, at 5% or better; ***, at 1% or better. *t*-Statistics are listed in parentheses below coefficient estimates. Y_t is total household wage and transfer income. UE_{t-1} is the unemployment expectations index. The instruments are the same as the second set used in Carroll, Fuhrer, and Wilcox (1994). The balance sheet variables are the growth in total household liabilities ($\Delta \log D_{t-1}$), the debt service burden (rD_{t-1}/Y_{t-1}), and the ratio of total household liabilities to annuity income (D_{t-1}/A_{t-1}). θ is the estimated coefficient on the moving average error term. A constant term was also included but is not reported.

tion of nondurables, our measure of consumption is spending on nondurable goods from the NIPA accounts.[8]

Results are contained in Table 1. Our first regression reproduces the basic result of Campbell and Mankiw (1989): the coefficient on predictable income growth is enormously statistically significant (with a *t*-statistic greater than 4), and suggests that rule-of-thumb consumers earn

three-month T-bill rate, the change in the unemployment rate, and the growth of the S&P 500 index; and one lag of the log difference between consumption and income and of the measure of sentiment being tested (in our case, unemployment expectations; in Carroll, Fuhrer, and Wilcox's paper, overall consumer sentiment). The adjusted R^2 on the first-stage regression for income growth is 0.41.

8. The model is often estimated on the sum of nondurables and services consumption. However, in the final version of NIPA data, substantial parts of services consumption are constructed using quarterly interpolation through annual estimates, where the later endpoint for the interpolation is strictly in the future of some of the quarterly estimates it is used to construct. This potentially introduces spurious time-series properties into the services component of spending, which are most easily avoided by excluding services from the measure of consumption. For more discussion of these points, see Wilcox (1992).

roughly half of aggregate labor income. Our second regression performs a simple Hall-style test of whether lagged unemployment expectations are useful in predicting current consumption growth. Again the answer is overwhelmingly yes; the *t*-statistic is 3.7. Our next regression reconfirms the main result of Carroll, Fuhrer, and Wilcox (1994): the lagged level of consumer sentiment (as measured by unemployment expectations) contains substantial predictive power for consumption growth *even after controlling for the information sentiment contains about income growth.*[9,10]

Turning now to the role of balance-sheet variables, our goal is to test whether such variables add anything to the benchmark sentiment-augmented Campbell–Mankiw model presented in row 3 of Table 1. In our background empirical work we examined a broad set of measures of household balance-sheet conditions, but in this paper we present results for only three measures: the ratio of liabilities to annuity labor income, the ratio of liabilities to assets, and the growth rate of liabilities.[11] None of the other balance-sheet variables we examined performed better (in the sense of being more highly correlated with the dependent variables we are interested in) than these three variables.[12]

Our empirical test is simply whether lagged balance-sheet variables are statistically significant when we add them to the sentiment-augmented Campbell–Mankiw model.[13] As rows 4 through 6 of the table show, none of the balance-sheet variables is statistically significant in any of the

9. Carroll, Fuhrer, and Wilcox used the overall index of consumer sentiment rather than the unemployment expectations index we use here; also, they tested for the joint significance of four lags of sentiment, rather than just a single lag as we do.
10. When lagged unemployment expectations are added to the Campbell–Mankiw equation, the coefficient estimate on forecastable income growth is about half of its previous value and just misses being statistically significant (the *p*-value is 0.103). The reason the statistical significance of the forecastable part of income growth drops so dramatically when lagged unemployment expectations are included in the regression is that lagged unemployment expectations are highly correlated with the forecastable component of income growth. Whether income growth is significant, lagged unemployment expectations are significant, or neither is significant is somewhat sensitive to the choice of instruments; in particular, if the instrument set does not contain variables that provide substantial information about income growth that is independent of the information about income growth contained in unemployment expectations, typically neither income growth nor unemployment expectations is individually significant.
11. See below for a discussion of how we constructed our estimate of annuity labor income.
12. We also examined the ratio of debt to net worth, the ratio of debt to liquid assets, the ratio of debt to current income, and the ratio of the debt service burden to annuity income, among others.
13. Of course, we also add them to the set of instruments used for predicting income growth.

regressions.[14] Thus, there is little evidence that household balance-sheet conditions have any influence on nondurables consumption growth that operates through any channel outside of the sentiment-augmented Campbell–Mankiw model.[15]

We now turn to the question of the relative importance for nondurables consumption of *innovations* to annuity income and to unemployment expectations. This question is of central importance to the enterprise of this paper, because the answer should help to inform us whether ignoring fluctuations in uncertainty is a small omission (which is well worth the associated modeling dividend of analytical tractability) or a large omission, so that any model which ignores uncertainty is likely to tell a seriously incomplete story about the determinants of consumption over the business cycle.

To examine this issue (and many others we will introduce later in the paper) we need an estimate of the level of annuity income. We construct two estimates, first following a method used to estimate annuity personal disposable income in the FRB-US model at the Federal Reserve Board, then using a method of our own devising. The FRB-US methodology ($A_t^{\text{FRB-US}}$) is based on an assumption that the ratio of personal income to GDP is stationary and that the GDP gap is stationary. A VAR forecasting system is used to estimate the projected future output gap XGAP and the projected future gap in the ratio of income to GDP, YGAP. The VAR system includes equations for inflation, the fed funds rate, XGAP,

14. The debt-to-annuity-income variable appears to be nonstationary, while consumption growth is approximately stationary; econometric theory implies that for a large enough time sample, the coefficient in a regression of a stationary variable on a nonstationary one must yield a zero coefficient, so the insignificance of this variable is hardly surprising.

15. These results are somewhat at variance with previous results of Ludvigson (1996), who found that predictable debt growth was significantly related to consumption growth. We were able to reproduce Ludvigson's results, and have determined that there are four reasons for the differences in outcomes. First, our measure of consumption spending is restricted to nondurable goods, while Ludvigson followed most of the previous literature by examining spending on nondurable goods and services. We believe that the data construction methods for the quarterly services expenditures render those data unsuitable for regressions of this kind. Second, because our focus is on the overall structure of household balance sheets, our measure of debt is total household liabilities, while Ludvigson's balance-sheet variable was consumer installment credit, i.e., mainly debt exclusive of mortgages. Third, Ludvigson's test was whether consumption growth was related to predictable debt growth, while our test is a more direct test of the Campbell–Mankiw model: whether lagged debt growth matters. Finally, Ludvigson was using the standard Campbell–Mankiw model as her baseline rather than the sentiment-augmented model we are using [although our result that lagged debt growth is insignificant holds up even when we estimate a standard (non-sentiment-augmented) Campbell–Mankiw model].

and YGAP. We also added four lags of income growth and the unemployment expectations variable to each equation.[16]

Our own measure of annuity labor income (A_t^{Ours}) is created by forecasting the present discounted value of the sum of the next two years of labor income using a set of forecasting variables drawn from Carroll, Fuhrer, and Wilcox's (1994) set of instruments for income growth. We make the assumption that beyond two years income is expected to grow at a constant rate equal to the average growth rate over the entire sample period. Using this growth rate, we calculate the annuity value of income from two years to infinity and add this to the forecasted discounted sum of income over the next two years to get A_t^{Ours}. For more details on the two methods of constructing annuity income, see the companion methodology paper Carroll and Dunn (1997).

In principle, if our estimate of the innovation to annuity income were perfect (or, more realistically, if the variables we use to construct the measure were valid instruments for annuity income growth) then the following equation would characterize nondurable consumption growth in the Campbell–Mankiw model:

$$\Delta \log C_t = (1 - \lambda)\, E_{t-1}\rho^{-1}(r_t - \delta) + \lambda\, \Delta \log Y_t + (1 - \lambda)\, \Delta \log A_t. \qquad (1)$$

Hence we could obtain an estimate of the fraction of income accruing to rule-of-thumb consumers from the coefficient on *actual* current income growth in a regression of consumption growth on current income growth and the current innovations to annuity income.[17] Table 2 presents the results when equation (1) is estimated using our two measures of annuity income.

The first regression shows that the *lagged level* of UE and the *current innovation* to our measure of annuity income are roughly equally important in explaining current consumption growth. The second regression shows that when the current innovation to UE is added to the equation, neither it nor the innovation to annuity income is individually statisti-

16. We are grateful to David Reifschneider at the Federal Reserve for explaining the FRB-US methodology to us. Because we are adapting the FRB-US methodology to a purpose quite different from its intended purpose, and because we are using a different measure of income, any empirical inadequacies of the annuity income measure we construct using the FRB-US methodology should be laid at our doorstep, not the FRB-US model staff's.

17. This point relies heavily on the assumption that our estimate of annuity income growth correctly captures all the implications for annuity income of the innovation to current income. However, we do include current income growth among the variables used to construct annuity income, so in principle any such information is indeed included.

Table 2 EFFECTS OF INNOVATIONS ON NONDURABLES CONSUMPTION[a]

Row	$\Delta \log Y_t$	$\Delta \log A_t^{\text{Ours}}$	$\Delta \log A_t^{\text{FRB-US}}$	UE_{t-1}	ΔUE_t	$\bar{R}^2$	D–W
1	0.326	0.186		−0.833		0.33	1.83
	(3.15)***	(2.82)***		−(2.55)***			
2	0.324	0.124		−1.003	−0.907	0.34	1.92
	(3.15)***	(1.59)		−(2.92)***	−(1.52)		
3	0.391		0.189	−0.654		0.29	1.92
	(3.41)***		(1.20)	−(2.00)**			
4	0.394		0.000	−0.981	−1.413	0.32	2.00
	(3.50)***		(0.00)	−(2.83)***	−(2.47)**		

[a]Dependent variable is nondurable consumption growth, quarterly data, 1963:3–1994:3. *, significant at 10% or better; **, at 5% or better; ***, at 1% or better. *t*-Statistics are listed in parentheses below coefficient estimates. Standard errors were constructed using a serial-correlation-robust covariance matrix (allowing serial correlation at lags up to 8). Y_t is total household wage and transfer income. A_t is annuity labor income. UE_{t-1} is the unemployment expectations index. A constant term was also included but is not reported.

cally significant; however, the lagged level of UE remains important. The last two regressions show that, after controlling for unemployment expectations, the FRB-US measure of annuity income provides no further information about consumption growth at all.

In sum, the standard model of nondurable consumption growth, the Campbell–Mankiw model, implies that consumption growth should be related to two variables: income growth and the innovations to annuity income. Our empirical work shows that unemployment expectations are at least as important as either of these traditional variables in explaining nondurables consumption growth. Lagged balance-sheet variables, on the other hand, are essentially uncorrelated with nondurable consumption growth once unemployment expectations are controlled for.

2.2 BALANCE SHEETS AND SPENDING ON DURABLE GOODS AND HOUSING

The standard CEQ PIH model described above applies to consumption of nondurable goods and services. However, as Mankiw (1982) showed, the model can be expanded to provide implications about durable-goods spending if sufficient assumptions are made. In particular, if there are no transactions costs associated with durable-goods purchases and if durable goods enter the utility function in a Cobb–Douglas manner, it is

Table 3 CONSUMPTION OF DURABLES, BASELINE EQUATION[a]

Annuity income measure	A_{t-1}/A_t	Prime_t	UE_{t-1}	ΔUE_t	Y_t/A_t	$\bar{R}^2$	*D–W*
A_t^{Ours}	−0.213	−0.115	−2.326	0.702	0.219	0.44	0.55
	−(3.22)***	−(3.16)***	−(6.11)***	(1.03)	(2.80)***		
$A_t^{\text{FRB-US}}$	0.329	−0.136	−2.931	−1.246	0.328	0.75	0.83
	(2.65)***	−(4.97)***	−(9.35)***	−(2.07)**	(10.40)***		
$A_t = Y_t$	−0.368	−0.104	−1.809	0.475	0.058	0.52	0.56
	−(3.24)***	−(2.71)***	−(3.73)***	(0.65)	(0.73)		

[a]Dependent variable is the ratio of durables consumption to annuity labor income, 1963:3–1994:3. *, significant at 10% or better; **, at 5% or better; ***, at 1% or better. *t*-Statistics are listed in parentheses below coefficient estimates. Standard errors were constructed using a serial-correlation-robust covariance matrix (allowing serial correlation at lags up to 18). Prime_t is the prime rate. Y_t is total household wage and transfer income, and A_t is annuity labor income. UE_{t-1} is the unemployment expectations index. The balance-sheet variables are the growth in total household liabilities ($\Delta \log D_{t-1}$), the debt service burden (rD_{t-1}/Y_{t-1}), and the ratio of total household liabilities to annuity income (D_{t-1}/A_{t-1}). Household net worth, the ratio of current income to annuity income, and a constant term were also included as independent variables but are not reported.

possible to show that the ratio of the stock of durable goods Z_t to annuity income A_t should be constant[18]:

$$Z_t = \omega A_t. \tag{2}$$

Expenditure on durable goods in this case will be determined by two factors: the spending needed to counteract depreciation, and the spending required to adjust the stock of durable goods to any changes in the level of annuity income:

$$E_t^z = Z_t - (1 - \delta)Z_{t-1}, \tag{3}$$

$$E_t^z/A_t = \omega - (1 - \delta)\omega A_{t-1}/A_t. \tag{4}$$

Table 3 presents empirical results when we estimate an equation like (4) using U.S. NIPA data on durables expenditures, augmented with

18. The assumption of frictionless adjustment is of course unattractive for durable goods, as many authors have pointed out. For an excellent discussion of the literature and of the difficulties, see Bertola and Caballero (1990), who also propose a sophisticated (and complicated) method of estimating the process for durables expenditures under a generalized (*S*,*s*) model with fixed return points. See also Bertola and Caballero (1994) and Eberly (1997). For reasons that will become clear in the theoretical discussion below, however, these frameworks are not well suited to addressing the issues we are interested in here of the relationship between labor-income uncertainty, balance-sheet variables, and spending. We therefore adopt the approach of estimating as simple an empirical model as possible, with an eye to finding any correlations sufficiently robust that any theoretical model should be consistent with them.

UE_{t-1} and ΔUE_t. We also include: the ratio of current income to annuity income, to allow some scope for current income to affect spending directly; the prime rate, to allow a channel for interest rates; and the ratio of net worth to annuity income (not shown in the table, to save space; it was usually not statistically significant). We present results separately for our estimate of annuity income, the annuity income estimate based on the FRB-US methodology, and the analogous results where we use current income rather than an estimate of annuity income.[19] We experimented with several methods of removing low-frequency movements or trends in the data, but they had little effect and are therefore not included.[20]

When the measure of annuity income is A^{Ours}, the annuity income ratio A_{t-1}/A_t gets the correct (negative) sign (implying that strong growth in annuity income from $t-1$ to t is associated with high durables purchases), as does the interest rate $Prime_t$. However, the lagged level of unemployment expectations is much more statistically significant than either annuity income or interest rates. Once again, the change in unemployment expectations does not enter significantly. Finally, the ratio of current income to annuity income, which plays no role in determining durables spending in the CEQ PIH model, is also highly significant in our regressions. This result differs from Bernanke (1984), who found in household data that transitory shocks to income had no effect on durables purchases. The discrepancy suggests either that our annuity income measures are imperfect or that consumers do in fact buy durables when they receive windfalls.

The second row of the table presents results when annuity income is measured using the FRB-US methodology. The main difference in results is that the annuity-income ratio now receives the wrong sign. The last row of the table shows the results when current income, rather than an estimate of annuity income, is used as a divisor. Results are generally similar to those for our measure of annuity income.

The top panel of Table 4 shows the results when our balance-sheet variables are added to the baseline durables regression.[21] The debt-to-

19. For the Y_t/A_t variable, we use the ratio of current income to our estimate of annuity income.

20. The Durbin–Watson statistics in the table indicate a large amount of positive serial correlation in durables spending. Mankiw (1982) shows that in the model we use the level of spending should follow a white-noise process, and so the empirical finding of severe serial correlation is inconsistent with the model. Caballero (1993) shows, however, that an (S,s) model implies precisely such slow adjustment. Because our theoretical model is essentially an expanded (S,s) model, Caballero's (1993) logic should apply to our model as well.

21. For brevity, we report only the results for A^{Ours}. Conclusions are similar for $A^{FRB\text{-}US}$.

Table 4 CONSUMPTION OF DURABLES AND LAGGED BALANCE-SHEET VARIABLES[a]

Row/Measure	A_{t-1}/A_t	Prime_t	UE_{t-1}	ΔUE_t	Y_t/A_t	*Balance-sheet measure*	$\bar{R}^2$	D–W
Entire sample period (1963:3–1994:3)								
1 $\Delta \log D_{t-1}$	−0.185	−0.095	−1.131	0.790	0.150	0.377	0.54	0.85
	−(3.13)***	−(2.95)***	−(2.45)**	(1.27)	(2.13)**	(4.22)***		
2 rD_{t-1}/Y_{t-1}	−0.217	−0.103	−2.906	0.497	0.183	0.413	0.50	0.65
	−(3.22)***	−(3.54)***	−(6.97)***	(0.79)	(2.27)**	(2.94)***		
3 D_{t-1}/A_{t-1}	−0.220	−0.115	−2.229	0.415	0.299	−0.027	0.48	0.57
	−(3.46)***	−(3.20)***	−(6.57)***	(0.64)	(5.13)***	−(2.64)***		
Before financial liberalization (1963:3–1980:1)								
4 $\Delta \log D_{t-1}$	−0.196	−0.007	−2.025	−0.407	0.236	0.180	0.79	1.77
	−(4.22)***	−(0.31)	−(7.87)***	−(0.95)	(7.79)***	(3.91)***		
5 rD_{t-1}/Y_{t-1}	−0.189	−0.017	−2.527	−0.682	0.273	0.010	0.75	1.53
	−(3.52)***	−(0.74)	−(10.10)***	−(1.40)	(8.63)***	(0.06)		
6 D_{t-1}/A_{t-1}	−0.143	−0.106	−2.098	−0.670	0.275	0.057	0.78	1.65
	−(2.62)***	−(2.02)**	−(6.97)***	−(1.45)	(9.32)***	(2.43)**		

[a]Dependent variable is the ratio of durables consumption to annuity labor income. *, significant at 10% or better; **, at 5% or better; ***, at 1% or better. *t*-Statistics are listed in parentheses below coefficient estimates. Standard errors were constructed using a serial-correlation-robust covariance matrix (allowing serial correlation at lags up to 18). Prime_t is the prime rate. Y_t is total household wage and transfer income, and A_t is our measure of annuity labor income. UE_{t-1} is the unemployment expectations index. The balance-sheet variables are the growth in total household liabilities ($\Delta \log D_{t-1}$), the debt service burden (rD_{t-1}/Y_{t-1}), and the ratio of total household liabilities to annuity income (D_{t-1}/A_{t-1}). Household net worth, the ratio of current income to annuity income, and a constant term were also included as independent variables but are not reported.

annuity-income ratio again gets a negative and significant coefficient using our measure of annuity income. However, both the lagged debt growth and the lagged debt service burden are positive and significant. Note that this is the opposite of what would be expected if precarious balance-sheet conditions tended to deter consumers from spending. Instead, the regressions indicate that consumers tend to spend more on durable goods during periods when the debt service burden has been high or recent debt growth has been high. The obvious interpretation is that these results reflect a simultaneity problem: factors that cause consumers to be willing to spend heavily on durable goods also tend to make them willing to tolerate high debt service burdens or rapid debt growth or high ratios of debt to assets.

One specific hypothesis is that the simultaneity problem reflects the financial liberalization of the 1980s, which may have allowed consumers to borrow more in order to purchase durable goods. If this explanation is correct, the statistical significance of the relationship between the durables spending ratio and balance-sheet variables should have been much weaker in the period before financial liberalization. The bottom panel of the table therefore presents results for the same sets of regressions, but restricting the sample to the period before 1980. Evidence for the debt service burden is consistent with the liberalization hypothesis: it is insignificant during the earlier time period. The results for lagged debt growth also lend some support to the idea; although the variable remains highly statistically significant, the coefficient estimates for the pre-1980 period are about half of their values over the entire period. Finally, the debt-to-annuity-income ratio now receives a *positive* and significant coefficient.

We now briefly examine the evidence on spending on what Saddam Hussein might call the mother of all durable goods: housing. Table 5 presents regressions patterned on our durable-goods regressions, but where the dependent variable is the number of homes sold per capita and the interest rate is the average rate on new mortgages.[22] For the baseline regression specification, the results are remarkably similar (given the totally independent sources of data) to those for durables spending: Coefficient estimates on every variable are betweeen two and four times the coefficient estimates in the durables regression, and the patterns of statistical significance are also very similar. Results for the balance-sheet variables are also similar to those for the durables regressions, though more exaggerated, in that both lagged debt growth and

22. To save space in the table, we do not report the coefficient on a trend variable, which was highly statistically significant in all regressions. We obtained similar results with alternative methods of detrending. We also report results only for our measure of annuity income.

Table 5 TOTAL HOME SALES[a]

Row/Measure	A_{t-1}/A_t	Mort_t	UE_{t-1}	ΔUE_t	Y_t/A_t	*Balance-sheet measure*	$\bar{R}^2$	D–W
1	−0.929	−0.698	−7.471	−1.541	1.172		0.51	0.33
	−(3.48)***	−(4.82)***	−(4.21)***	−(0.70)	(2.99)***			
2 $\Delta \log D_{t-1}$	−0.681	−0.600	−2.341	−1.721	0.784	1.306	0.62	0.85
	−(2.79)***	−(4.82)***	−(1.27)	−(0.77)	(2.54)***	(3.78)***		
3 rD_{t-1}/Y_{t-1}	−0.896	−0.499	−8.962	−2.834	1.226	0.920	0.51	0.34
	−(3.21)***	−(2.23)**	−(5.08)***	−(1.56)	(3.20)***	(1.26)		
4 D_{t-1}/A_{t-1}	−0.709	−0.600	−8.679	−4.295	1.206	0.205	0.58	0.42
	−(2.50)**	−(4.84)***	−(4.66)***	−(2.38)**	(3.34)***	(2.85)***		

[a]Dependent variable is total home sales per capita, 1972:3–1990:1. Annuity income constructed using our method. *, significant at 10% or better; **, at 5% or better; *** at 1% or better. *t*-Statistics are listed in parentheses below coefficient estimates. Standard errors were constructed using a serial-correlation-robust coveriance matrix (allowing serial correlation at lags up to 18). The measure of home sales is new and existing single-family home sales per capita. Mort_t is the effective rate on conventional home mortgage loans. Y_t is total household wage and transfer income, and $A_t$is annuity labor income. UE_{t-1} is the unemployment expectations index. The balance-sheet variables are the growth in total household liabilities ($\Delta \log D_{t-1}$), the debt service burden (rD_{t-1}/Y_{t-1}), and the ratio of total household liabilities to annuity income (D_{t-1}/A_{t-1}). Household net worth, a constant term, and a 9-year centered moving average of home sales were also included as independent variables but are not reported.

lagged debt service burden receive coefficients more than four times as large as in the durables regressions. However, the lagged debt-to-annuity-income ratio, which received a negative and significant coefficient in our baseline durables regressions, is positive and significant here.

Our conclusion is that spending on durables and housing is very robustly correlated with lagged unemployment expectations. It is also highly correlated with our measure of annuity-income growth, and with the ratio of current income to annuity income. However, with the exception of debt growth, durables spending is not robustly correlated with any balance-sheet measure we examined.[23] Given the enormous changes in the U.S. financial system over the period our data covers, and given the endogenous nature of balance-sheet positions, it is perhaps not surprising that most balance-sheet measures do not bear any stable relationship to spending. Indeed, the surprise may be that one balance-sheet measure, debt growth, *does* seem to bear a relatively stable relationship to spending. We therefore turn now to an exploration of the determinants of debt growth.

2.3 THE CYCLICAL DYNAMICS OF DEBT GROWTH

Aside from the sharp increase in the debt ratio beginning in the mid-1980s, perhaps the most interesting feature of our Figure 1 was that debt appears to exhibit a distinct cyclical pattern: its growth rate is much slower during recessions (the shaded regions of the chart) than during expansions.

It is a bit difficult to pin down the representative-agent CEQ PIH model's implications for debt, because the model does not distinguish debt from assets; aggregate net worth and human wealth are sufficient statistics for aggregate behavior. Of course, the vast majority of debt is associated with purchases of homes and other durable goods, so to the extent that our earlier empirical work captures the dynamics of home sales and durables purchases, the remaining interesting question to ask about debt growth is what else it is correlated with. The way we answer this question empirically is to see what variables are statistically significant explanators of debt growth once we control for contemporaneous home sales. The results are presented in Table 6.

As usual, the first variable we examine is lagged unemployment expectations; as usual, it is highly statistically significant and negative. Debt growth is also negatively correlated with the change in unemployment

23. This conclusion is consistent with recent work by Garner (1996), who found that most measures of the household debt burden do not Granger-cause durable-goods expenditures or GDP, and McCarthy (1997), who finds in a VAR framework that debt measures have little effect on subsequent nondurable- or durable-goods spending.

Table 6 DETERMINANTS OF DEBT GROWTH[a]

Row/Measure	H_t	UE_{t-1}	ΔUE_t	$\Delta \log A_t^{Ours}$	*Balance-sheet measure*	θ	SSR	D–W
1	0.196 (4.64)***					0.539 (5.85)***	0.59	2.46
2	0.140 (5.79)***	−2.169 −(5.72)***				0.244 (2.15)**	0.55	2.15
3	0.131 (5.78)***	−2.864 −(6.34)***	−1.970 −(3.90)***			0.306 (2.72)***	0.49	2.21
4	0.133 (6.35)***	−2.536 −(7.38)***		0.180 (3.90)***		0.202 (1.69)*	0.51	2.12
5	0.130 (5.90)***	−2.867 −(6.41)***	−1.662 −(2.42)**	0.059 (0.79)		0.287 (2.51)***	0.49	2.19
6 $\Delta \log D_{t-1}$	0.045 (2.98)***	−1.385 −(5.25)***			0.588 (7.84)***	−0.443 −(6.42)***	0.48	2.07
7 rD_{t-1}/Y_{t-1}	0.133 (6.12)***	−2.345 −(6.09)***			0.063 (0.82)	0.218 (1.86)*	0.54	2.13
8 D_{t-1}/A_{t-1}	0.147 (5.85)***	−2.063 −(5.17)***			−0.004 −(0.60)	0.259 (2.24)**	0.54	2.17

[a]Dependent variable is $\Delta \log D_t$, growth in total household liabilities: quarterly data, 1968:2–1994:3. *, significant at 10% or better; **, at 5% or better; ***, at 1% or better. *t*-Statistics are listed in parentheses below coefficient estimates. H_t is home sales per capita and A_t is annuity income. UE_{t-1} is the unemployment expectations index. The balance-sheet variables are the lagged dependent variable ($\Delta \log D_{t-1}$), the debt service burden (rD_{t-1}/Y_{t-1}), and the ratio of total household liabilities to annuity income (D_{t-1}/A_{t-1}). θ is the estimated coefficient on the moving-average error term. A constant term was also included but is not reported.

expectations, although (as usual) at a much lower level of statistical significance than the correlation with the lagged level. Again, a possible interpretation is that the statistical significance of these variables owes to some correlation they have with the level of future income; but, as in all our previous regressions, when a measure of the change in annuity income is added to the equation, the statistical significance of lagged unemployment expectations is unaffected (although the annuity-income growth variable is also significant). Finally, debt growth is uncorrelated with the lagged values of our other two balance-sheet variables but is significantly positively autocorrelated.

These regressions suggest that there is an independent channel for unemployment expectations in influencing debt growth, even beyond whatever effects unemployment expectations have on home sales. Because we found earlier that the pace of home sales is itself negatively influenced by unemployment expectations, in a sense these results imply that unemployment expectations are doubly important for debt growth.

Implicit in our entire discussion up to this point has been an assumption that the pattern of debt over the business cycle is determined by consumers' unconstrained choices. An alternative possibility is that debt growth slows over the business cycle not because consumers desire to borrow less but because lenders restrict credit. A large literature now exists suggesting that lenders tighten credit standards to businesses during recessions, so that only high-quality borrowers are able to borrow freely in bad times; see Bernanke, Gertler, and Gilchrist (1996) for a survey. A recent paper by Bernanke, Ferri, and Simon (1997) presents evidence from the Federal Reserve's *Survey of Consumer Finances* suggesting that a similar phenomenon may afflict consumers.

One way to identify demand and supply effects is to examine a form of mortgages for which there should be no cyclical effect on supply. The best candidate here is mortgages issued by the Veterans' Administration, because by law these mortgages are available to all qualified former military personnel. Because the government assumes the default risk, the supply of this form of mortgage financing should not fluctuate over the cycle even if lenders become more risk-averse in recessions. Indeed, because the government bears the risk of VA mortgages, one would expect to see a relative *increase* in the supply of VA mortgages. If the supply of other forms of credit does decline, we would also expect to see an increase in the relative demand for VA mortgages; hence any declines in VA mortgage issuance over the cycle probably *underestimate* the pure demand effect.

Figure 3 plots the number of VA mortgages originated in each quarter since 1981, together with total mortgages originated over the same period. There is clearly a strong correlation between VA mortgages and non-VA

Figure 3 VA MORTGAGE ORIGINATIONS AND TOTAL MORTGAGE ORIGINATIONS OVER TREND GDP

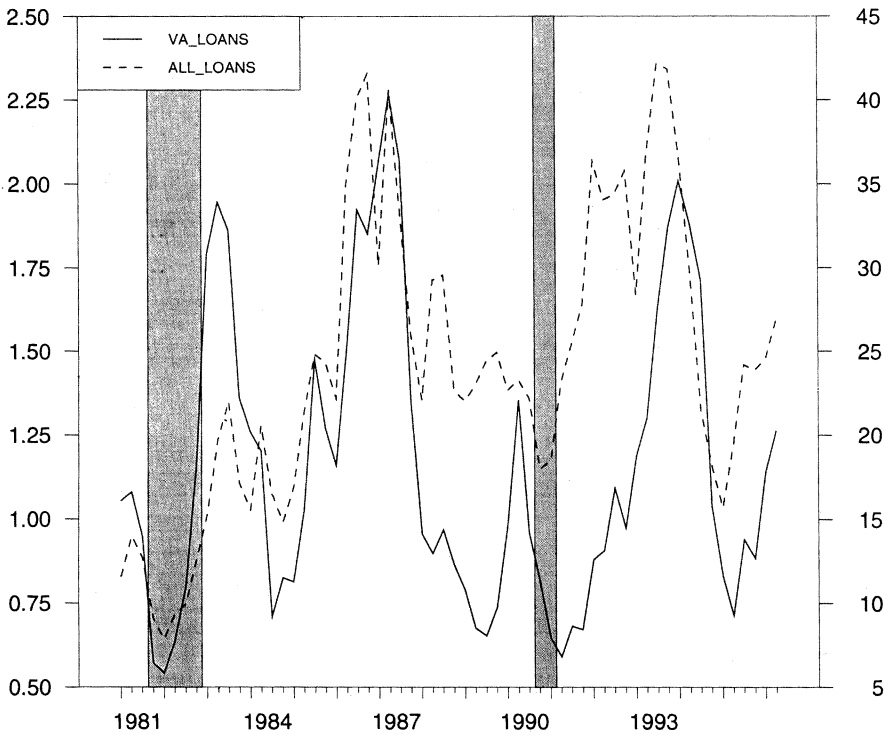

mortgages. Furthermore, during the two recessions in the sample, VA mortgages appear to fall, if anything, by more than non-VA mortgages. This evidence strongly suggests that demand factors play a very important role in fluctuations in mortgage borrowing over the business cycle.

This completes our discussion of the cyclical characteristics of consumption spending, home sales, and household balance sheets. We draw several conclusions. First, spending for nondurables, durables, and housing all generally responds to changes in annuity income (or at least our measure of annuity income) in the direction implied by the frictionless CEQ PIH model, although the magnitude of the response is generally not nearly so large as the model would predict. Second, unemployment expectations typically seem to play at least as important a role as changes in annuity income in determining spending decisions. However, most of the information content of unemployment expectations variables is captured by the *lagged level* of unemployment expectations

rather than by the change in unemployment expectations. Finally, the only measure of household balance-sheet positions that is robustly correlated with spending is the lagged growth rate of debt.

We turn now to the question of whether a model which incorporates a serious treatment of uncertainty, transactions costs, and liquid assets can explain the broad pattern of our empirical results.

3. *The Model*

3.1 THEORY

The consumer's objective is to maximize expected discounted utility from consumption of housing services Z and nonhousing goods C. The period utility function is CRRA in a Cobb–Douglas aggregate of utility from nonhousing consumption and the stock of housing:

$$u(C_t, Z_t) = \frac{(C_t^{1-\alpha} Z_t^{\alpha})^{1-\rho}}{1-\rho}, \tag{5}$$

where $\rho = 2$ is the coefficient of relative risk aversion.

There are five state variables which constrain or influence the consumer's choice of C and Z: the current stock of spendable resources X_t [the sum of wealth and current labor income Y_t; or "cash-on-hand" in Deaton's (1991) terminology], the size of the home (if any) the consumer owns at the beginning of the period H_t^b, the level of the consumer's permanent labor income P_t, an indicator I_t for the aggregate state of the economy, and the consumer's current employment status J_t (for *job*). Note that we do not list mortgage debt as one of the state variables. This is because we make sufficient assumptions to guarantee that the ratio of the mortgage debt to home value is constant, thereby reducing the number of state variables in the problem by one. The critical assumption is that the mortgage payment in each period contains a term that corresponds to the depreciation rate of the home. Hence the balance owed on the mortgage shrinks in each period by the same fraction that the value of the home shrinks.

The consumer's choices within each period are determined as follows (and as summarized in the table below). First the consumer makes a homeownership decision. If the consumer begins the period owning no house ($H_t^b = 0$), the decision is whether or not to buy a house, whose value we will denote $H_t^e = \phi P_t$, i.e., we assume that consumers must by a house whose value is equal to $\phi = 3$ times their real after-tax permanent income, in accord with standard rules of thumb in the housing industry, (see Fannie Mae Foundation, 1997). Buyers must also put up a down payment of amount $d = 0.2$ of the value of the house, and pay fees and

taxes in amount $b = 0.03$. Renters purchase housing services in optimally chosen amount Z_t at price $q\lambda$, where λ is the flow cost of homeownership,[24] and the restriction $q = 1.5 > 1$ gives consumers an incentive to buy. If the consumer begins the period as a homeowner, they can sell the house and rent (implying $H_t^e = 0$), keep the house they currently own ($H_t^e = H_t^b$), or sell the current house and buy a new one. For homeowners, the flow of housing services is equal to the size of the house, $Z_t = H_t^e$.

Given our assumption that debt depreciates at the same rate as the house, the outstanding amount of debt will always be given by the amount $(1 - d)H_e^t$. We assume that this debt must be serviced in each period by a fixed mortgage payment $m = \delta + r$, where $r = 0.02$ is the after-tax real rate of return and $\delta = 0.02$ is the depreciation rate of the house. The presence of the term δ in the mortgage payment represents the lender's compensation for the erosion in the real value of debt (this term can be thought of as roughly reflecting inflation).

Denoting the level of liquid assets that the consumer ends the period holding S_t, we can summarize the foregoing possibilities in the following table:

Initial status	*Period-t action(s)*	S_t *(end-of-period savings)*	H_t^e	Z_t
$H_t^b = 0$	Keep renting	$X_t - C_t - q\lambda Z_t$	0	Optimal
$H_t^b = 0$	Buy	$X_t - C_t - (d + b)H_t^e - [m(1 - d) + n]H_t^e$	ϕP_t	H_t^e
$H_t^b > 0$	See and rent	$X_t - C_t + (d - b)H_t^b - q\lambda Z_t$	0	Optimal
$H_t^b > 0$	Hold	$X_t - C_t - [m(1 - d) + n]H_t^e$	H_t^b	H_t^e
$H_t^b > 0$	Sell and buy	$X_t - C_t + (d - b)H_t^b - (d + b)H_t^e$ $-[m(1 - d) + n]H_t^e$	ϕP_t	H_t^e

We are now in position to write down the consumer's optimization problem. The consumer of course has no influence over the aggregate state; furthermore, for simplicity we assume that the consumer's job status and permanent income also evolve exogenously. Hence the control variables potentially available to the consumer are three: C, H^e, and Z. The Bellman equation is therefore

$$V_t(X_t, H_t^b, I_t, J_t, P_t) = \max_{\{C_t, Z_t, H_t^e\}} u(C_t, Z_t) + \beta E_t V_{t+1}(X_{t+1}, H_{t+1}^b, I_{t+1}, J_{t+1}, P_{t+1}),$$

where all notation is standard and $\beta = 0.96$ is the time-discount factor.

24. Equal to the lost interest on the capital tied up in the house plus depreciation costs plus maintenance costs.

The level of permanent labor income is assumed to follow a first-order Markov process with drift parameter G_{t+1}:

$$P_{t+1} = G_{t+1}P_t\Pi_{t+1}, \tag{6}$$

where Π_{t+1} is a stochastic shock to permanent labor income, and G_{t+1} is the mean growth rate for the permanent income of employed consumers given the aggregate state that prevails in period $t + 1$. With this process for permanent labor income, along with the fact that the utility function is homogeneous of degree zero,[25] it is possible to rewrite the problem in terms of ratios of C, Z, X, and H^b to permanent labor income, thus effectively reducing the number of state variables to four. Specifically, defining $c_t = C_t/P_t$ and z_t, x_t, and h_t^b similarly, the problem can be written as

$$v_t(x_t,h_t^b,I_t,J_t) = \max_{\{c_t,z_t,h_t^e\}} u(c_t,z_t) + \beta E_t(G_{t+1}\Pi_{t+1})^{1-\rho}v_{t+1}(x_{t+1},h_{t+1}^b,I_{t+1},J_{t+1}). \tag{7}$$

We assume that the level of actual labor income in period t is given by the level of permanent labor income multiplied by a transitory shock Ψ_t:

$$Y_t = P_t\Psi_t. \tag{8}$$

The consumer's decisions within the period determine the size of the housing stock at the end of the period H_t^e and the amount of liquid assets (or savings) on hand at the end of the period S_t subject to a liquidity constraint that requires $S_t \geq 0$. Given H_t^e and S_t, the levels of beginning-of-period housing H^b and cash on hand in period $t + 1$ are given by

$$H_{t+1}^b = (1 - \delta)H_t^e,$$
$$X_{t+1} = RS_t + Y_{t+1},$$

where $R = 1.02$ is the annual gross interest rate between periods. Dividing both sides of both of these equations by P_{t+1} and substituting from the permanent-labor-income equation (6) yields

$$h_t^b = \frac{h_t^e(1 - \delta)}{G_{t+1}\Pi_{t-1}},$$
$$x_{t+1} = \frac{R}{G_{t+1}\Pi_{t+1}}s_t + \Psi_{t+1}.$$

25. Plus certain conditions that must be (and are) satisfied by the constraints.

3.2 THE AGGREGATE STATE

Following the work of Sichel (1993, 1994), we assume that the aggregate economy has three states: recessions, which are characterized by high job loss risk and low aggregate growth; booms, which are characterized by low job loss risk and high aggregate growth: and recoveries, which always follow recessions and which exhibit high growth but continuing high job loss risk.[26] Transitions between these states are governed by the following Markov transition matrix:

		Period-t + 1 aggregate state		
		Expansion	*Contraction*	*Recovery*
Period-*t* aggregate state	Expansion	0.95	0.05	0
	Contraction	0.05	0.70	0.25
	Recovery	0.25	0.05	0.70

where the switching probabilities were chosen to match the empirical fraction of the time the economy has spent in expansion vs. contraction in the postwar United States, and the probabilities for the recovery period were chosen so that recoveries would last for four quarters on average and so that the probability of slipping from recovery back into recession would be the same as the probability of entering a recession from an expansion.

3.3 THE HOUSEHOLD INCOME PROCESS

3.3.1 The Employment State Unemployment spells last one or two periods, and when consumers lose their jobs, they know whether the spell will be a one- or a two-period spell (we chose this structure to allow the average spell length to be longer during recessions than during expansions). Consumers in the last period of an unemployment spell face the same employment hazards as employed consumers; thus a very unlucky consumer could experience two (or even more) unemployment spells in a row. Designating the status employed as E, unemployed with one remaining quarter of unemployment as U_1, and unemployed with two remaining quarters as U_2, we assume the employment state transition matrix in expansions is

26. The "recovery" phase allows our model to capture the fact that the unemployment rate typically remains higher than average for an extended period after the NBER trough.

		Period-t + 1 status		
		E	U_1	U_2
Period-t status	E	0.97	0.01	0.02
	U_1	0.97	0.01	0.02
	U_2	0	1	0

while we assume that in contractions and recoveries the matrix is

		Period-t + 1 status		
		E	U_1	U_2
Period-t status	E	0.96	0	0.04
	U_1	0.96	0	0.04
	U_2	0	1	0

where the transition probabilities were chosen to generate steady-state unemployment rates around 5% in expansions and 8% in contractions and recoveries (by "steady-state" we mean the rate that would eventually prevail if the economy remained in the expansion, or contraction, or recovery for many periods).

3.3.2 The Transitory Shocks Transitory shocks to income are drawn for all employed consumers in each period from a three-point symmetric distribution with mean one and equal probability mass on each of the three possible draws. Thus the possible draws are $(1 - \nu^e, 1, 1 + \nu^e)$ in expansions and $(1 - \nu^{cr}, 1, 1 + \nu^{cr})$ in contractions and recoveries, $\nu^{cr} \geq \nu^e$ (in practice we assume transitory shocks are of equal size in all aggregate states, $\nu^{cr} = \nu^e = 0.1$). Unemployed consumers receive unemployment compensation in amount YP_t with certainty, where we assume that the replacement rate $Y = 0.5$ does not vary with the cycle.

3.4 THE PERMANENT SHOCKS

For employed consumers, permanent shocks to income, like transitory shocks, are drawn in each quarter from a three-point symmetric distribution with mean one and equal probability mass on each of the three possible draws. We assume the three possibilities are (0.95, 1.00, 1.05) in all three aggregate states, which amounts to a conservative estimate, given that microeconomic studies typically estimate that the standard deviation of the annual innovation to permanent income is at least 10% (see Carroll, 1992, for a brief survey). We assume that unemployment spells in all three states of the economy typically end with consumers

taking jobs at a level of permanent income that is on average 10% lower than the permanent income associated with their previous job (this is one of the few statistics we were able to calibrate using existing data from the labor economics literature; see, e.g., Carrington, 1993, for evidence on the typical size of wage losses). However, we were unable to find evidence on how this statistic varies over the business cycle, so we assume that it is the same in all three aggregate states. We again assume a three-point symmetric distribution with equal probability weights on all three outcomes, but we assume that the shock process during contractions and recoveries is a mean-preserving spread of the shock process during expansions. Specifically, the possible outcomes are (0.8, 0.9, 1.0) in booms and (0.7, 0.9, 1.1) in contractions and recoveries.

3.5 SUMMARY

Although the model can be solved for quite general combinations of parameter values, we have intentionally kept the structure of uncertainty simple in order to make the model easier to understand and analyze. In our parametrization, the only differences in risk between aggregate states come from the fact that in recessions and recoveries unemployment spells are more likely, last longer, and are associated with larger permanent income shocks. The processes for transitory and permanent shocks for employed consumers are the same in all three aggregate states, as is the mean of the distribution for permanent shocks for the unemployed. Many of these parameters could in principle be calibrated using microeconomic data, but we were not able to find many existing studies that were useful for that purpose.

3.6 A WISH LIST

In order to solve the model, we had to make a variety of simplifying assumptions. Even so, the full version of the model used for analysis of the effects of financial market deregulation has six state variables: the four described above (x_t, h_t^b, I_t, J_t) plus the current value of the down-payment ratio d required for new-home purchases and the value of the down-payment ratio that prevailed when the consumer took out their mortgage loan. The full model takes our new Unix workstation four days to solve and another two to simulate, so substantially relaxing the simplifying assumptions is not feasible with present technology. It is nevertheless worthwhile to draw attention to the assumptions we would most like to relax as technology advances. First is the assumption that the level of debt is perfectly correlated with the level of the housing stock. We would have preferred to make assumptions that guaranteed at least a modest buildup of home equity over the course of time. The second assumption we would

like to relax is that there is no house price risk. Although Fratantoni (1996) found that the effects of this kind of risk were small compared to the effective risk caused by the fixed mortgage commitment, it would be useful to see whether that result carries over into this context. This assumption could obviously interact with the first assumption, because house price risk could put some consumers "under water," holding a mortgage whose value exceeds that of the house. Finally, we would like to allow consumers to choose the size of the new house they buy. However, we suspect that this last change would not affect behavior much: because consumers will live in their house for an average of ten years, it seems unlikely that transitory factors such as the current aggregate state should optimally have much effect on the optimal size of house to buy.

3.7 SOLUTION

As anyone familiar with the recent literature on consumption under uncertainty would anticipate, solution of this model was a major challenge. A short companion paper (Carroll and Dunn, 1997) briefly describes our solution method, which involves numerical iteration on the value function. Carroll and Kimball (1996) have shown that even in the simpler case where there is only a single, nondurable consumption good, the consumption policy rule is strictly concave (and therefore presumably not analytically soluble) whenever utility is of the hyperbolic absolute risk aversion (HARA) form (a class that subsumes constant absolute risk aversion (CARA), constant relative risk aversion (CRRA), and Stone–Geary versions of CARA and CRRA utility) and there is both labor-income and rate-of-return risk. That paper shows that there are only three degenerate cases which yield linear consumption rules: quadratic utility, CARA utility with only labor-income risk, and CRRA with only rate-of-return risk. Given the lack of analytical solutions to even the simpler problem for nondurable consumption, the resort to numerical methods was inescapable here—even if the fixed transaction costs did not add further major complications.

Previous work on (S,s) models has either assumed assumed risk neutrality of consumers (Bertola and Caballero, 1990) or that the only risk consumers face is rate-of-return risk (Grossman and Laroque, 1990; Eberly, 1997) in order to exploit the linearity of the optimal consumption rule under power utility [which, under certain further assumptions, implies a closed-form solution to even the more complicated (S,s) problem]. A very recent paper by Caplin and Leahy (1997) makes substantial progress in deriving empirical implications of a model in which the marginal utility of wealth does not vary over the business cycle (except as a result of interest-rate fluctuations). While these assumptions are

defensible for many purposes, they are obviously unacceptable in a study of the effects of labor-income uncertainty on durables purchases.

Despite the mathematical difficulty of solving the model, the behavior of consumers in it can be described reasonably simply. Most of the time they are homeowners, because ownership is cheaper than renting. During most of the time that they are homeowners, they engage in *buffer-stock saving*, in which they try to maintain a target level of liquid precautionary assets, which they use to smooth nonhousing consumption in the face of income shocks (see Deaton, 1991, and Carroll, 1992, 1997, for detailed analysis of buffer-stock saving behavior in a model with only nondurable goods). As the time approaches to buy a new home, however, they engage in a bit of extra saving in order to accumulate the required down payment.

The homeownership decision can be described as following a modified (S,s) rule. Because the value of the house depreciates over time, and because permanent labor income grows, the ratio of home value to permanent labor income drifts down over time. When this ratio drops far enough, the consumer sells the existing home and buys a new one. The most important twist in this model, relative to the standard (S,s) model of durable goods, is that the precise trigger point at which the consumer decides to buy a new house depends on both the anticipated risk of unemployment and the size of the consumer's current buffer stock of liquid assets. This is illustrated by Figure 4, which shows the lower

Figure 4 THE JUMP IN THE LOWER (S,s) TRIGGER

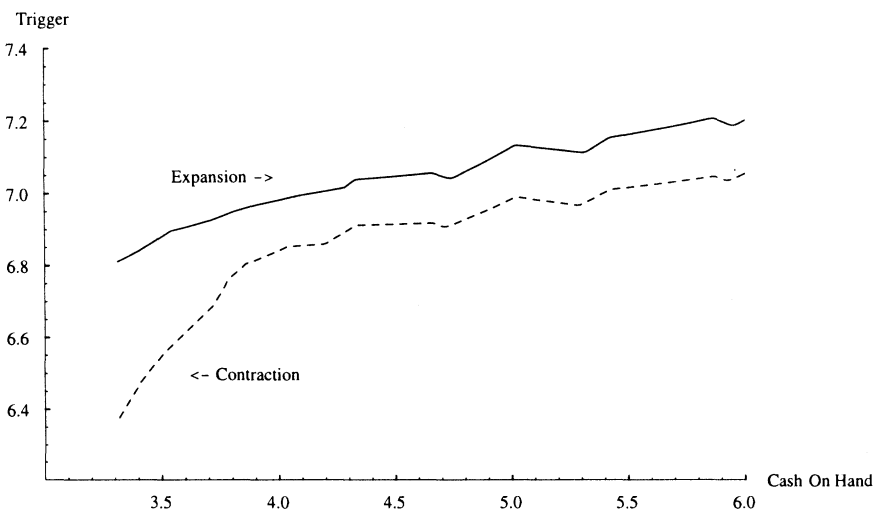

trigger point of the (*S*,*s*) rule as a function of the level of liquid assets the consumer has, for an employed consumer living in an economy in an expansion.[27] The curve is upward sloping, indicating that consumers with more liquid assets will buy a new durable earlier (or, more formally, at a higher trigger value). Note that, in the presence of aggregate shocks to transitory income, this result could rationalize our empirical finding that durables and home sales are high in periods of high transitory income. That is, when they receive windfall income, some consumers are pushed *rightward* across the (S,s) barrier. This is an interesting theoretical difference with the CEQ model as explored, for example, by Bernanke (1985), in which transitory shocks to income should have essentially no effect on durable-goods spending.[28]

The figure also shows (the dashed line) how the trigger locus changes if the economy enters a recession: for any given level of liquid assets, the trigger point is lower (consumers will put up with living in a poorer house rather than buy). That is, a consumer who had been on the brink of home purchase before the economy entered the recession will now wait until the house has depreciated more before buying. Alternatively, a consumer with a given house value will require a larger stock of precautionary liquid assets before he will be willing to buy. This shift in the lower (*S*,*s*) trigger is what we refer to in the title of the paper as "Jumping (S,s) Triggers."

The foregoing story is somewhat different from the standard (S,s) model's explanation of durables purchases over the business cycle found in, for example, Bar-Ilan and Blinder (1992), Bertola and Caballero (1990), or Caplin and Leahy (1997).[29] The main difference is the explicit importance of cyclical variation in labor-income uncertainty in our model: in the standard model, the sharp drop in durables purchases in recessions is triggered, not by an increase in uncertainty, but by a decrease in the *level* of expected future income and thus of "permanent income" as they define it. The empirical distinction between the two models is thus that our model would imply a strong effect of uncertainty

27. We also assume that the consumer bought his current house with an 80% mortgage and expects to finance the new house with an 80% mortgage.
28. One way to think about this finding is as an increase in the marginal propensity to consume durable goods out of transitory income. As a theoretical matter, this result corresponds closely to Kimball's (1990) finding that precautionary saving boosts the marginal propensity to consume nondurables out of transitory income.
29. One interesting recent paper that adopts a rather different approach to these issues is Greenspan and Cohen (1997), who model vehicle sales as a function of "scrappage" and who make a distinction between "engineering scrappage" and "cyclical scrappage." Roughly speaking, however, it is possible to interpret the effects of the jumping (*S*,*s*) trigger in our model as corresponding to the "cyclical scrappage" term in the Greenspan–Cohen model.

per se on durables purchases, even after controlling for permanent (or annuity) income. Another way to interpret the jump in the trigger is as reflecting the fact that an increase in uncertainty causes an increase in the marginal utility of liquid wealth, because its value as a buffer stock against uncertainty rises. This is in explicit contrast with Caplin and Leahy's assumption that the marginal utility of wealth is constant.[30]

For purposes of cyclical anlaysis, the most important implication of the model comes from the interaction of the precautionary saving motive and the jumping (S,s) bands. When the economy switches into a recession, a large proportion of the entire set of consumers who had been on the brink of home purchase suddenly feel that their current stock of precautionary saving, which had been adequate when they anticipated continued prosperity, is inadequate in the new, riskier environment. These consumers postpone their home purchases until they have accumulated enough additional precautionary savings to again feel comfortable with the home purchase decision (or until their home has deteriorated so much that they are willing to risk buying a new one even with a low buffer stock of liquid assets).[31]

Another interesting feature of this model that is not present in the standard model is that home equity serves as an additional reserve of emergency precautionary resources beyond liquid assets. Consumers who experience a particularly vicious series of income shocks can, in the last resort, sell their houses in order to tap the equity to finance current consumption. Of course, they pay a heavy price for this: They must incur brokerage fees and pay for rented housing services at a price substantially higher than the user cost of ownership. Still, extreme circumstances call for extreme measures. This feature of the model is interesting because several papers in the empirical literature on precautionary saving have found larger effects of uncertainty on net worth than on liquid assets. Carroll and Samwick (1997) speculate that the reason may be precisely this potential use of home equity as a precautionary reserve.

30. One recent paper which focuses on the effects of jumping (S,s) triggers is by Adda and Cooper (1997), who examine the effects of two natural experiments thoughtfully provided to economists by the French government. The experiments involved subsidies to automobile scrappage, which should have had the effect of moving the lower (S,s) trigger up. Adda and Cooper document that the reaction of automobile sales to the tax subsidies was quite similar to the predictions of an (S,s) model when the lower trigger moves up.

31. In the Greenspan–Cohen model, the implication would be that "cyclical scrappage" is strongly related to unemployment expectations. Although Greenspan and Cohen do not report regressions of cyclical scrappage on unemployment expectations, they do report that cyclical scrappage falls when the unemployment rate rises, which is roughly what one would expect from our model.

Our paper is not the first to argue that variations in the degree of uncertainty are important in explaining durables purchases over the business cycle. As Bernenke (1983) pointed out, and many authors have emphasized since, an increase in uncertainty increases the *option value* of waiting until the uncertainty is resolved.[32] A formal illustration of this can be seen in Eberly (1997); she shows that in a model with only rate-of-return risk, when the degree of rate-of-return risk goes up the (S,s) bands widen, provoking a response similar to that of the jump in the (S,s) band we depict. However, the underlying cause of the jump is rather different. In Eberly's model the primary reason for the shift in the (S,s) bands is that if the bands did not change, an increase in uncertainty would increase the expected present discounted value of the adjustment costs the consumer would have to pay. Thus the effect of uncertainty in her model has little to do with precautionary behavior—instead, it mainly reflects a change in the trade-off between minimizing average fixed costs and minimizing average distance from the optimal housing stock. Again, a useful way to understand the difference between the models is to realize that the main effect driving the jump in the (S,s) trigger in our model is an increase in the marginal utility of liquid assets—an effect that is absent in the Bernanke and Eberly models.

An even earlier analysis of many of these issues can be found in three insightful articles by Frederick Mishkin (1976, 1977, 1978) which anticipate many (though not all) of the theoretical results that come from our formal optimizing model. In particular, Mishkin (1978) argues that "A consumer suffering financial distress, and unable to pay his bills readily, would prefer holding highly liquid financial assets. This implies that as the consumer perceives an increasing probability of financial distress, he will decrease his demand for consumer durables and limit his purchases." Using an intuitive but ad hoc functional form, Mishkin also documents a strong correlation between durables purchases and consumer sentiment, and explicitly interprets consumer sentiment as a measure of uncertainty.

4. *Simulation Results: A Stylized Business Cycle*

Our simulation results examine the aggregate characteristics of an economy populated by 20,000 consumers behaving according to the optimal decision rules that solve the maximization problem in Section 3. As preparation for the simulations, we start the model economy off at an

32. For a thorough and recent treatment, see Dixit and Pindyck (1994).

essentially arbitrary point, then simulate for 400 quarters of expansion, by which time it has settled into a stochastic steady state with a reasonably stable distribution of consumers across the state space.

The first experiment we perform is to examine a recession of typical length (4 quarters) followed by recovery period of the same length. We show the path of aggregate variables from 8 quarters before the beginning of the recession to 4 quarters after the end of the recovery. Results are shown in Figures 5 and 6; the contraction is shaded dark gray, and the recovery period is shaded light gray. In the first quarter of the recession, the unemployment rate begins moving up as the new, higher job loss risk affects its first batch of victims. Recall, however, that unemployment spells in recessions last 2 quarters: this means that the new reces-

Figure 5 A TYPICAL RECESSION IN OUR SIMULATED ECONOMY

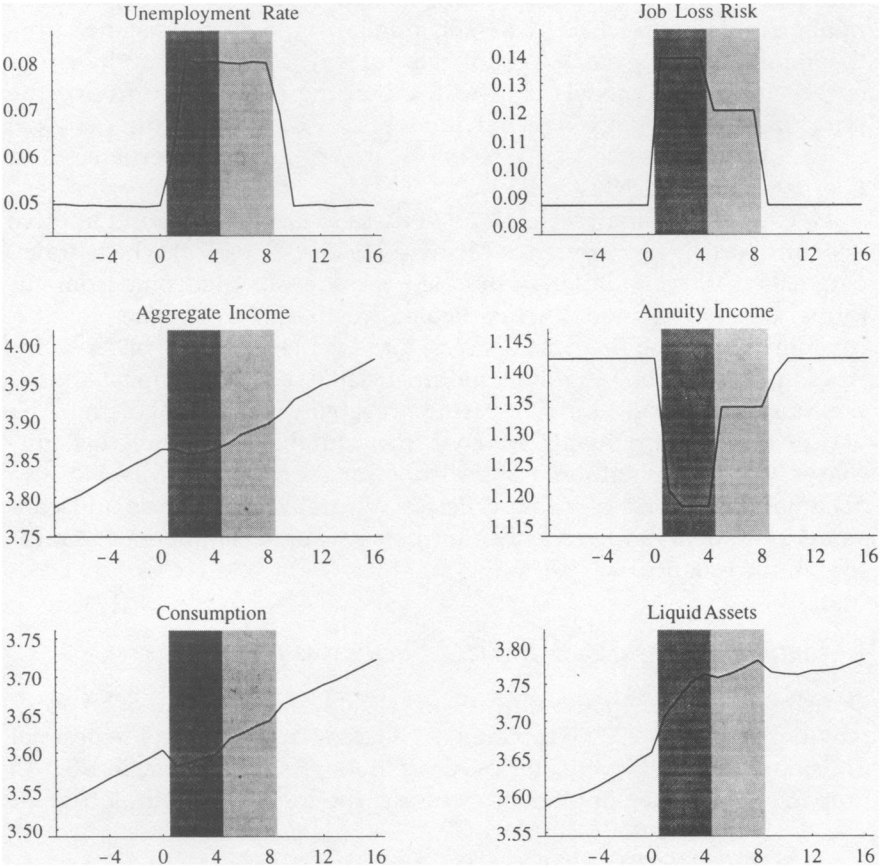

Figure 6 THE HOUSING MARKET AND BALANCE SHEETS

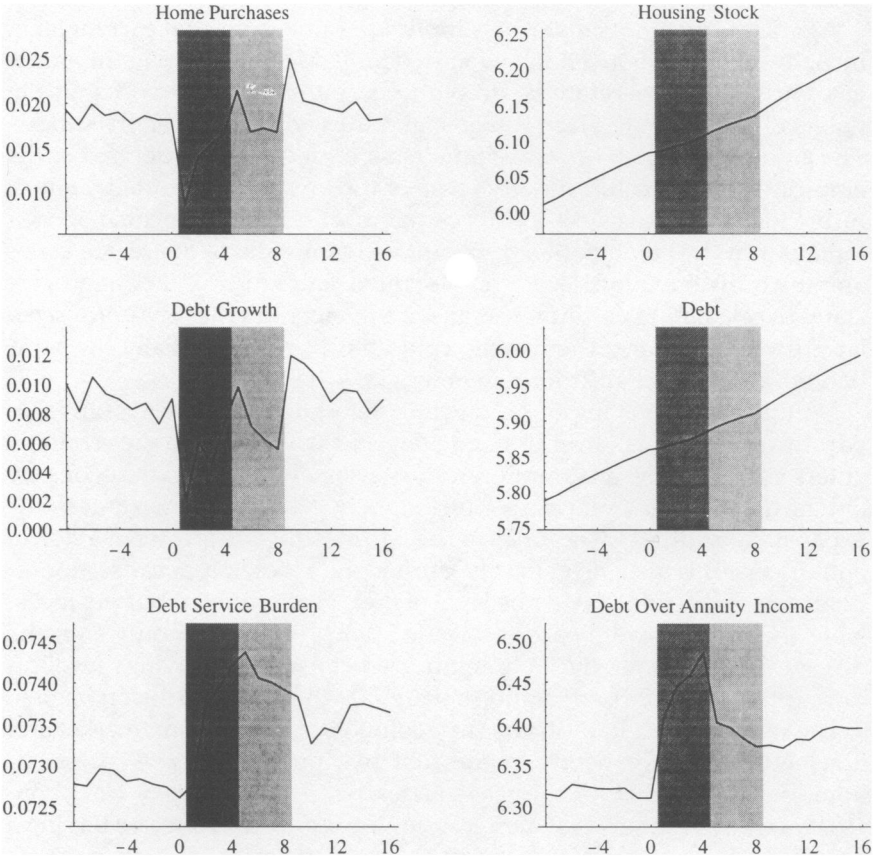

sionary equilibrium level of the unemployment rate is only reached in the second quarter of recession. Thereafter the unemployment rate stays at the same high level throughout the recession and recovery periods, reverting to its expansionary level only in the second quarter after the end of the recovery period.

The adjacent chart shows the expected risk of job loss over the next 4 quarters for a currently employed consumer, a statistic we take to be the closest analogue in our model to the unemployment-expectations variable used in our empirical work. Because it is an expectational variable, when the economy enters a recession this measure jumps immediately to its recessionary steady-state value. When the economy moves into the recovery phase, the expected job loss risk drops because consumers

know that it is likely that the economy will enter an expansionary phase in the near future.

Aggregate income is given by simply summing the actual current labor income of all the households we are simulating. Movements in aggregate income can therefore be decomposed into those due to changes in the level of permanent labor income and those due to transitory shocks. The adjoining figure shows annuity income, calculated as defined in the empirical section as the annuity value of the present discounted value of future labor income. We calculate aggregate annuity income in our model from the combination of the transition matrix for aggregate states and the transition matrices for employment states during each aggregate state. In calculating annuity income we abstract from the long-term secular growth in income; the results would have been essentially the same had we allowed the drift term to enter.

Nondurables consumption, which is determined (as always) in large part by expectations, drops immediately and sharply when the economy enters a recession. Consumption recovers somewhat when the economy enters the recovery phase and further when the economy enters the expansionary phase. The final figure in the set shows the behavior of liquid assets, which rise sharply during the recession because households feel the need to boost the level of their precautionary buffer stocks. Note that the precautionary motive is intense enough to outweigh the dissaving being done by the unemployed consumers. Savings level off during the recovery period and remain flat when the expansion begins.

The next set of figures shows the evolution of the housing market and household balance sheets. In the first two quarters of the recession, home sales plummet for the reasons described above: newly wary consumers want a higher level of precautionary liquid assets before buying a house. Note the impressive magnitude of the initial decline in home sales: the rate of sales per capita falls by roughly 50%. After the initial collapse, home sales begin to rise again, then show a minor surge when the economy enters the recovery phase. Finally, when the economy switches into expansion there is a massive surge of home sales as the consumers who had been postponing purchases for precautionary reasons throw caution to the wind.

This last phenomenon, the surge of sales when the economy exits the contraction, has a natural interpretation as the release of "pent-up demand." "Pent-up demand" is a phrase used loosely by analysts of the housing and auto sectors who claim that recessions are periods when "pent-up demand" rises, only to be "released" when the economy emerges from the contraction. Pent-up demand could be defined rigor-

ously in our model as the demand which would be immediately be expressed in purchases if consumer expectations returned to normal levels. In other words, pent-up demand corresponds to the set of consumers populating the region of the (*S,s*) diagram between the jumping loci of the lower (S,s) band.[33]

It is worth noting just how close the correspondence is between this phenomenon in the formal model and the informal descriptions of industry analysts in both the housing and the automotive markets. For example, a May 11, 1992 editorial (p. 12) in *Automotive News* read, in part: "[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."

The graph to the right of the home-purchases graph in Figure 6 shows the obvious implication of purchases for the level of the housing stock: at the onset of a recession, the growth rate of the housing stock decelerates. Below are the growth rate of debt and the time path of the debt stock, which strongly resemble the patterns of home purchases and the housing stock.

The bottom two panels of Figure 6 show the behavior of our measures of household balance-sheet conditions over the business cycle. Both the debt service burden and the ratio of debt to annuity income rise sharply at the beginning of the recession, in both cases because the numerator is largely fixed by past decisions while the denominator (income or annuity income) falls when the economy enters a recession. Thereafter the debt service burden drifts up until the economy enters a full expansionary phase again, whereas the ratio of debt to annuity income drops as soon as the economy enters the recovery period (because the level of annuity income jumps up; see the previous set of figures).

5. *Comparing the Model with U.S. Cyclical Data*

We turn now to some simulations based on the pattern of expansion and contraction for the U.S. economy since 1961, roughly the period for which we were able to perform our empirical work on U.S. NIPA data. Again we start the economy off from the steady-state equilibrium achieved after 400 quarters of continuous expansion, but for quarters 401 through 539 (corresponding to 1962:2 through 1995:4) we set the aggre-

33. This definition differs somewhat from the definition proposed by Caballero and Engel (1994). They investigate a model with fixed (*S,s*) bands and describe a period of high pent-up demand as a period with a heavier than usual concentration of agents near the (unmoving) trigger point.

Figure 7 SIMULATED ECONOMY WITH ACTUAL RECESSION PATTERN

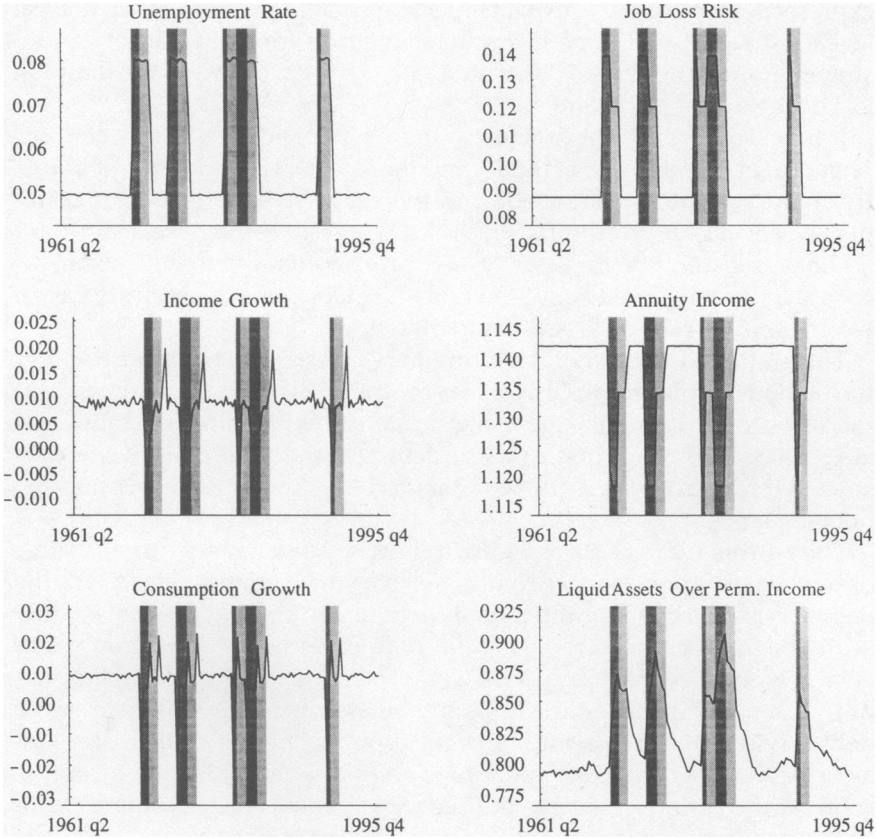

gate state of the simulated economy equal to the aggregate state of the corresponding quarter for the U.S. economy as indicated by the official NBER chronology. (We arbitrarily assume that every recession is followed by a recovery period that is 4 quarters long, which is the expected duration implied by the transition matrix.) Figures 7 and 8 shows the results graphically.

5.1 NONDURABLES CONSUMPTION GROWTH

We begin by examining the analogue to the Campbell–Mankiw equation estimated in the first part of the paper. The top panel of Table 7 reproduces the baseline sentiment-augmented Campbell–Mankiw equation from Table 1. Row 1 in the second panel of the table shows that when a standard Campbell–Mankiw equation is estimated on the simulated non-

Figure 8 HOUSING, DEBT, AND BALANCE SHEETS

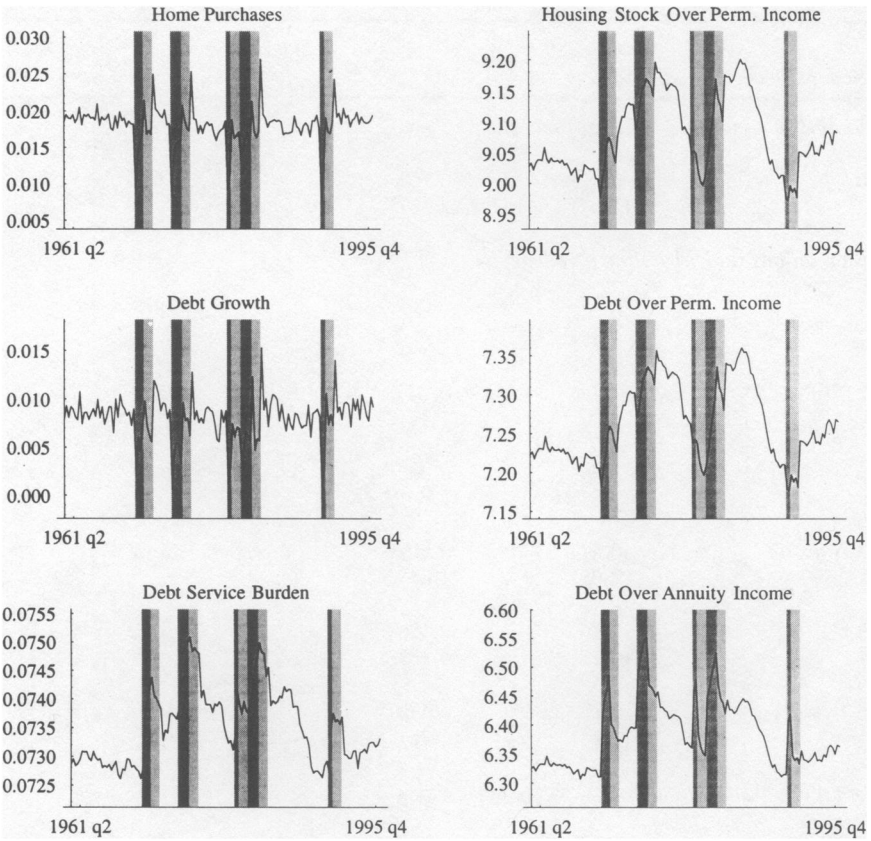

durables consumption data from our model, the forecastable component of income growth gets an insignificant and negative coefficient.[34]

Row 2 shows that consumption growth is not significantly related to lagged unemployment expectations, again in accord with the CEQ PIH model and at variance with the empirical results. Finally, when both predictable income growth and lagged unemployment expectations are included, neither is significant at the 5% level. Thus, under baseline parameter values the model does not reproduce the empirical results we found when estimating the Campbell–Mankiw model in Table 1.

34. Under some alternative assumptions on parameters, the model does reproduce the Campbell–Mankiw finding. Given how long it takes to solve the model, we were unable to explore the parameter space sufficiently to determine what kinds of parameter combinations generate the Campbell–Mankiw result.

Table 7 THE CAMPBELL–MANKIW MODEL ESTIMATED ON SIMULATED DATA[a]

Row/Measure	$E_{t-1}\,\Delta \log Y_t$	UE_{t-1}	*Balance sheet*	SSR	D–W
Empirical results (reproduced from Table 1)					
0	0.269 (1.64)	−0.906 −(2.18)**		0.50	1.98
Simulations under baseline parameter values					
1	0.032 (0.19)			0.55	1.93
2		0.404 (1.45)		0.55	1.81
3	0.244 (1.31)	0.633 (2.03)**		0.47	1.97
4 $\Delta \log D_{t-1}$	0.417 (1.92)*	0.507 (1.61)	−0.394 −(1.29)	0.42	1.98
5 rD_{t-1}/Y_{t-1}	0.182 (0.78)	0.459 (0.98)	0.359 (0.54)	0.49	1.97
6 D_{t-1}/A_{t-1}	0.229 (1.21)	0.495 (1.25)	0.007 (0.57)	0.47	1.97
Simulations after financial liberalization					
7	−0.127 −(0.62)			0.73	1.96
8		0.572 (1.87)*		0.67	1.93
9	0.135 (0.58)	0.698 (1.90)*		0.62	2.00

[a]Dependent variable is nondurable consumption growth. *, significant at 10% or better; **, at 5% or better; ***, at 1% or better. *t*-Statistics are listed in parentheses below coefficient estimates. Y_t is total household wage and transfer income. UE_{t-1} is the unemployment expectations index. The instruments are the same as those used in Carroll, Fuhrer, and Wilcox (1994). The balance-sheet variables are the growth in total household liabilities ($\Delta \log D_{t-1}$), the debt service burden (rD_{t-1}/Y_{t-1}), and the ratio of total household liabilities to annuity income (D_{t-1}/A_{t-1}). A constant term was also included but is not reported.

The next regressions examine the model's predictions for the explanatory power of lagged balance-sheet measures. In accord with our empirical results, none of the balance-sheet measures has any explanatory power for the growth of nondurables consumption.

Table 8 examines how nondurables consumption in our model responds to *innovations* to income and unemployment expectations; the corresponding U.S. empirical results from Table 2 are again reproduced in the top panel. Recall that the CEQ PIH model would imply a coefficient of 1 on $\Delta \log A_t$ and zero on all other variables, while the Campbell–Mankiw model with $\lambda = 0.5$ would imply coefficients of 0.5 on both $\Delta \log Y_t$ and $\Delta \log A_t$, but would still imply coefficients of zero on the UE

Table 8 EFFECT OF INNOVATIONS ON NONDURABLES CONSUMPTION GROWTH[a]

Row	$\Delta \log Y_t$	$\Delta \log A_t$	UE_{t-1}	ΔUE_t	$\bar{R}^2$	D–W
Empirical results (reproduced from Table 2)						
0	0.324 (3.15)***	0.124 (1.59)	−1.003 −(2.93)***	−0.907 −(1.52)	0.34	1.92
Simulations under baseline parameters						
1	0.109 (1.51)	1.323 (22.41)***	−3.398 −(2.41)**		0.86	2.44
2	0.078 (1.85)*		−4.470 −(5.43)***	−5.561 −(41.28)***	0.95	1.86
3	−0.006 −(.021)	0.486 (12.08)***	−5.860 −(10.07)***	−4.050 −(25.94)***	0.98	1.62
Simulations after financial liberalization						
4	−0.032 −(0.36)	1.530 (21.85)***	−4.394 −(2.67)***		0.86	2.50
5	−0.054 −(1.05)		−5.434 −(5.56)***	−63.926 −(39.55)***	0.95	1.76
6	−0.186 −(5.36)***	0.596 (13.75)***	−7.510 −(11.59)***	−46.023 −(27.59)***	0.98	1.57

[a]Dependent variable is simulated nondurable consumption growth. *, significant at 10% or better; **, at 5% or better; ***, at 1% or better. *t*-Statistics are listed in parentheses below coefficient estimates. Standard errors were constructed using a serial-correlation-robust covariance matrix (allowing serial correlation at lags up to 8). Y_t is total household wage and transfer income. A_t is annuity labor income. UE_{t-1} is the unemployment expectations index. A constant term was also included but is not reported.

variables. Row 3 in the second panel shows that under baseline parameter values our model implies a coefficient of about zero on $\Delta \log Y_t$ and 0.5 on $\Delta \log A_t$. However, the major difference between our model and either the CEQ PIH model or the Campbell–Mankiw model is our model's implication that both the lagged level and the change in UE should be highly statistically significant.[35] This constitutes at least a partial victory over the standard models, which provide no role at all for unemployment expectations *per se*. However, it is fair to say that the model is at best a modest success in explaining nondurables data, since it does not replicate the basic Campbell–Mankiw result. Furthermore, even for the unemployment expectations variable the match between theory and data is imperfect: the theory implies that the contemporaneous change in unemployment expectations should be vastly more important than the lagged level, but the empirical regressions found the opposite result. Carroll, Fuhrer, and Wilcox (1994) speculate that a model which incorporates both habit formation and labor-income uncertainty might be able to explain the importance of lagged uncertainty for current consumption growth; a recent paper by Overland (1997) provides a formal underpinning for this idea. Alternatively, it may take consumers time to formulate new spending plans upon receipt of new information; this could be formalized in a model in which consumers draw up budgets only periodically, and do not change their spending patterns until they find the time to draw up a new budget.

5.2 THE CYCLICAL DYNAMICS OF DURABLES SPENDING

Table 9 presents the results when we estimate equations for our simulated home sales data similar to those estimated earlier for both NIPA durable goods and total U.S. home sales; again the corresponding empirical result is reproduced in the first row of the table.[36]

In our simulated data the annuity-income ratio is insignificant, but both the lagged level of unemployment expectations and the change in unemployment expectations are highly significant. Here the level and

35. Because the variables are defined rather differently and scaled quite differently, it would not be appropriate to compare the coefficient estimates on UE from the model with those from the data; hence we examine only statistical significance.

36. Here we take the "corresponding" result from the table on durable-goods spending rather than the table on home sales. Although we calibrate our model to match certain features of the housing market, it is clear that under alternative parameter values the model could equally be interpreted as a model of purchases of automobiles or other durable goods. Given the similarity of the empirical results for home sales and durable-goods sales documented in Tables 3 and 5, it is of little consequence whether we compare our model's predictions with the pattern of durable-goods sales or home sales.

Table 9 HOME SALES IN THE SIMULATED ECONOMY[a]

Row/ Measure	A_{t-1}/A_t	UE_{t-1}	ΔUE_t	Y_t/A_t	*Balance-sheet measure*	$\bar{R}^2$	D–W
Empirical baseline (reproduced from Table 5)							
0	−0.542	−7.471	−1.541	1.172		0.51	0.33
	−(3.48)***	−(4.21)***	−(0.70)	(2.99)***			
Simulations under baseline parameters							
1	0.037	−5.260	−19.407	−0.106		0.76	1.80
	(0.85)	−(8.07)***	−(10.55)***	−(2.57)***			
2 $\Delta \log D_{t-1}$	0.048	−4.047	−20.452	−0.081	0.179	0.77	2.16
	(1.14)	−(5.50)***	−(11.31)***	−(1.98)**	(3.20)***		
3 rD_{t-1}/Y_{t-1}	0.043	−5.780	−19.641	−0.100	0.154	0.76	1.79
	(0.98)	−(7.63)***	−(10.66)***	−(2.41)**	(1.33)		
4 D_{t-1}/A_{t-1}	0.026	−4.875	−19.227	−0.096	−0.002	0.76	1.79
	(0.57)	−(5.49)***	−(10.31)***	−(2.16)**	−(0.64)		
Simulations after financial liberalization							
5	−0.189	−7.020	−21.149	−0.296		0.80	1.93
	−(2.71)***	−(6.67)***	−(6.74)***	−(4.56)***			

[a]Dependent variable is the number of home sales. *, significant at 10% or better; **, at 5% or better; ***, at 1% or better. *t*-Statistics are listed in parentheses below coefficient estimates. A_t is annuity labor income, and $Y_t$is total household wage and transfer income. UE_{t-1} is the unemployment expectations index. The balance-sheet variables are the growth in total household liabilities ($\Delta \log D_{t-1}$), the debt service burden (rD_{t-1}/Y_{t-1}), and the ratio of total household liabilities to annuity income (D_{t-1}/A_{t-1}). A constant term was also included but is not reported.

the change in the level of the unemployment expectations index are roughly equally statistically significant.

Turning to the balance-sheet variables, lagged debt growth receives a positive and significant coefficient; recall that it was the only balance-sheet variable that was robustly significant in the NIPA data. Although simultaneity seemed the most plausible interpretation for the empirical results, there was no obvious way to prove that simultaneity was the correct interpretation. Here the answer is clear: simultaneity is the culprit. Debt growth is acting as a summary statistic for all of those characteristics of the aggregate environment which are important in determining the pace of home sales but are not captured by the other observed

aggregate variables. For example, during the course of recessions, home sales and debt growth both plummet initially, but recover substantially over the succeeding few quarters (even while the economy remains in recession). The partial recovery in sales reflects a combination of the buildup of consumers' buffer stocks of precautionary savings and the continuing depreciation of their homes [moving some of them across even a lowered (S,s) trigger]. These changing circumstances are not captured by our observed aggregate variables, but they are at least partly captured by lagged debt growth. Hence lagged debt growth's statistical significance is entirely attributable to the fact that it is an endogenous variable responding to unobserved but important real determinants of home sales.

The remainder of Table 9 shows that the other two lagged balance-sheet variables are not systematically related to home sales (as they were not in the empirical data). The reason can be seen in Figure 8: the debt-to-income ratio and the debt service burden tend to be high during recessions because income is temporarily low, but also tend to be high in recoveries and early expansions, because upon recovery the pent-up demand built up during the recession is satisfied by a large number of home purchases and a consequent runup in aggregate debt. Hence both variables tend to be higher than average both during periods of particularly low sales (recessions) and particularly high sales (early recoveries and expansions).

The fundamental question these regressions are designed to address is whether our model performs better than the standard models in explaining our empirical findings in Tables 3,4 and 5. On the whole, the answer is yes. Our model implies a very important role for unemployment expectations beyond any correlation they may have with current or expected future levels of income. And it provides an interpretation for the finding that lagged debt growth is consistently positively related to current home sales, and that other balance-sheet measures are not consistently related to home sales. However, as in the nondurables regressions, the model implies a much stronger reaction to innovations in uncertainty than we observe empirically. We speculated above that habit formation might explain the sluggishness of nondurables consumption with respect to unemployment-expectations innovations; for durable goods, however, time-to-build or decision-lag considerations seem more plausible. This is especially so for housing decisions; a consumer who has gone to the trouble of househunting, lining up financing, negotiating, and bidding on a house is unlikely to back out at the last moment because of a sudden change in unemployment expectations. Similar but less forceful arguments apply for automobile purchases.

5.3 THE CYCLICAL DYNAMICS OF DEBT GROWTH

In the model, the primary determinant of debt growth is home sales. Indeed, since all debt is used for home purchases, and since the value of all homes purchased is in exactly the same proportion to the permanent labor income of the buyer, one might think that the model implies that data on home sales should explain 100% of the variation on debt growth. A glance at Figure 6 will confirm that the patterns of some sales and debt growth over the cycle are indeed quite similar. However, the model does provide several channels through which other variables influence aggregate debt growth. First, a small number of consumers who have experienced a particularly nasty series of shocks find themselves forced to sell their homes and rent temporarily in order to get access to the emergency reserve of precautionary resources represented by their home equity. Second, among the consumers who are currently renting, fewer will be willing to buy new homes when unemployment expectations are pessimistic. The number of consumers who are forced to sell and rent will obviously be on average related to the level and change of unemployment expectations. Finally, note that—because the (S,s) trigger jumps around—even though every purchase represents exactly the same amount of debt acquisition (relative to the permanent income of the buyer), every sale does not reflect the same amount of debt retired. Hence we should expect variables that affect the location (S,s) trigger to have an effect on debt growth.

Table 10 presents the results when we estimate regressions for debt growth like those estimated in Table 6 above. As expected (and as in the empirical data), debt growth is very closely related to home sales; when the pace of home sales is the only regressor, the $\overline{R}^2$ is 0.76. However, the next regression shows that the lagged level of the unemployment expectations index does provide additional explanatory power for debt growth (again corresponding to the empirical result). When we add the growth rate of annuity income to this baseline regression, the innovation to annuity income is not statistically significant, in contrast with the empirical regressions. In contrast to the results for durable and nondurable goods, the change in unemployment expectations is *not* statistically significant. Finally, we consider the lagged balance-sheet measures, all of which are negatively correlated with current debt growth. These results contrast with the empirical regressions, in which the lagged dependent variable received a strongly positive coefficient and the other balance-sheet measures were insignificant.

In sum, the model captures (almost by assumption) the strong empirical correlation between home sales and debt growth, but, in contrast

Table 10 DEBT GROWTH AND ITS CORRELATES[a]

Row/ Measure	H_t	UE_{t-1}	ΔUE_t	$\Delta \log A_t$	Balance-sheet measure	$\bar{R}^2$	D–W
Empirical baseline (reproduced from table 6)							
0	0.130 (5.90)***	−2.867 −(6.41)***	−1.662 −(2.42)**	0.059 −(0.79)			2.12
Simulations under baseline parameters							
1	0.826 (20.95)***					0.76	2.07
2	0.812 (20.88)***	−1.419 −(2.81)***				0.77	2.22
3	0.805 (10.36)***	−1.465 −(2.24)***	−0.203 −(0.11)			0.77	2.22
4	0.791 (13.54)***	−1.564 −(2.67)***		0.017 (0.49)		0.77	2.23
5	0.806 (10.34)***	−1.477 −(2.25)**	0.726 (0.30)	0.026 (0.56)		0.77	2.24
6 $\Delta \log D_{t-1}$	0.826 (21.43)***	−2.324 −(3.79)***			−0.126 −(2.51)***	0.78	1.94
7 rD_{t-1}/Y_{t-1}	0.825 (21.25)***	−0.673 −(1.11)			−0.223 −(2.17)***	0.78	2.20
8 D_{t-1}/A_{t-1}	0.811 (20.88)***	−0.916 −(1.41)			−0.003 −(1.23)	0.77	2.21

[a]Dependent variable is the growth in total household liabilities: simulated data. *, significant at 10% or better; **, at 5% or better; ***, at 1% or better. *t*-Statistics are listed in parentheses below coefficient estimates. H_t is home sales per capita. UE_{t-1} is the unemployment expectations index. The balance-sheet variables are the lagged dependent variable ($\Delta \log D_{t-1}$), the debt service burden (rD_{t-1}/Y_{t-1}), and the ratio of total household liabilities to annuity income (D_{t-1}/A_{t-1}). A constant term was also included but is not reported.

with a CEQ PIH model or a standard (*S*,*s*) model, it also provides an interpretation for the empirical finding that unemployment expectations are significantly related to debt growth. It does not, however, imply the observed empirical positive autocorrelation in debt growth after unemployment expectations have been controlled for.

5.4 SUMMARY

The analysis of this section has attempted to determine whether our model does a better job than standard models of explaining the empirical regularities relating nondurable consumption growth, durables purchases, balance-sheet variables, and unemployment expectations. The model is successful in that it implies an important role for unemployment expectations in addition to the expected level of future income. However, it also suggests that there is a paradox about the role of unemployment expectations: while the model implies that consumption growth, durables purchases, and debt acquisition should be strongly affected by changes in unemployment expectations, our empirical work found that the lagged level of expectations was always much more statistically important than the change in expectations.

6. *Was the 1990 Recession "Special"?*

To this point in the paper we have not directly addressed the question of whether the 1990 recession was "special" in any sense, although we motivated the paper by noting that common analyses of the 1990 recession attributed the unusual consumption weakness to "household debt overhang." In this section we examine first the theory and then the evidence.

6.1 THEORY

6.1.1 The Dynamic Response to Deregulation As briefly noted earlier, prior to the 1990 recession there was a rapid and considerable runup in the ratio of household debt to income (see Figure 1). The most plausible explanation is that this was the consequence of the wide-ranging deregulation of financial markets that took place in the late 1970s and early 1980s.

Capturing the full complexity of financial deregulation is obviously beyond the scope of the model introduced above. However, both before and after deregulation, home mortgage borrowing was by far the largest component of total household debt. To the extent that the main effect of deregulation was to make mortgage borrowing easier by reducing required down payments, our model can be used to get a sense of the likely effects of deregulation. The particular experiment we consider is a one-time reduction in the down-payment requirement from our 20% baseline assumption to 10%. Of course, the progress of credit liberalization was in reality much more gradual, but this experiment should at least give a sense of the likely results of a more gradual deregulation.

Figure 9 DYNAMIC PATH OF ECONOMY AFTER FINANCIAL LIBERALIZATION

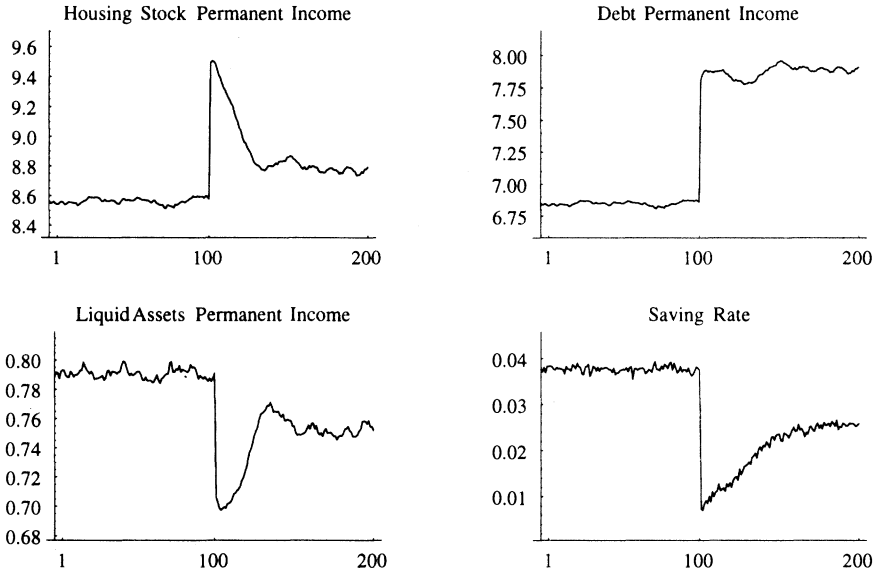

In the short term, the effects of deregulation are very similar to those of moving from a recession to an expansion: the bottom of the (S,s) band jumps upward instantly. Figure 9 depicts the results of reducing the down-payment requirement for an economy which had previously been in stochastic steady-state equilibrium. The immediate effect of deregulation is to spur an avalanche of home sales, which is accompanied by a massive runup in debt and consequently a large increase in the aggregate debt-to-income ratio. Eventually the level of housing per capita falls most of the way (although not all the way) back to its original steady-state level, but the ratio of debt to income plateaus at a substantially higher level. The level of liquid assets immediately drops sharply, as most of the consumers who had been saving up for a down payment now find that, in combination with the equity from their previous home, their current stock of liquid assets is enough to cover the new lower down-payment requirement. The level of liquid assets gradually rebounds a bit as new homeowners struggle to build up their buffer stocks of liquid assets to the target level, but the new steady-state level of liquid assets is well below its pre-deregulation equilibrium. This reflects the fact that a substantial part of the average stock of liquid assets represented saving for down payments rather than precautionary saving. The

consequences of deregulation for the aggregate saving rate are particularly interesting: in the three or four years after deregulation it drops from about 2¾% to under 1%, but eventually recovers a bit to settle down at slightly less than 2%. The U shape in the saving rate reflects the fact that for quite a while after deregulation most of the housing stock still consists of homes bought in the pre-deregulation period when the down-payment requirement was higher. These homeowners on average need to do very little down-payment saving, because the comparatively large equity in their previous home is by itself almost enough for the down payment on the new home. Eventually, however, the entire housing stock is composed of homes bought after liberalization and consumers have to boost their saving somewhat in order to accumulate down payments again.

6.1.2 Cyclical Properties of the Deregulated Economy From the standpoint of cyclical analysis, perhaps the most interesting question to ask about the deregulated economy is whether the higher prevailing debt burdens make aggregate consumption more volatile and in particular more responsive to unemployment expectations. We address this question by repeating the simulation and regression analysis of Section 5 for a deregulated economy that is otherwise identical to our baseline economy.

The bottom panels of Tables 7 and 8 present the results for the nondurables regressions in the deregulated economy. Results are on the whole not much different: both forecastable income growth and lagged unemployment expectations remain statistically insignificant, as do the balance-sheet variables (not reported). Nondurables consumption does react a bit more strongly to a change in unemployment expectations, but the change is modest.

The bottom panel of Table 9, however, shows that home sales are more sensitive to unemployment expectations in the high-debt economy: the coefficient on the lagged level of unemployment expectations changes from about -5 to about -7, and the coefficient on the change in unemployment expectations increases from about -19 to about -21.[37] Meanwhile, the annuity-income ratio (which was insignificant in the baseline economy) becomes statistically significant.

There are several reasons why home sales are more sensitive to uncertainty in the liberalized economy. The most important is probably simply that buying a house is a considerably riskier financial venture, for two reasons. First, and most important, there is a great deal less home equity available as an emergency reserve against major disasters (a long unem-

37. This increase in the significance of UE is the smallest increase we found under any combination of parameter values we checked. In the original draft of the paper, the coefficient on UE_{t-1} almost doubled.

ployment spell or a substantial reduction in the level of permanent income). Second is a mechanism emphasized by Fratantoni (1996): Because mortgage payments cannot be altered once the mortgage is taken out, all adjustment of consumption to any income shocks must be borne entirely by nondurables consumption. The larger mortgage payment associated with a lower down payment thus implies that at any given amount of liquid wealth, any given amount of uncertainty will have a greater influence on nondurables consumption.

One way to think about these results is to consider the large downpayment requirement as a form of forced saving which, essentially as a side effect, also serves a precautionary role. When the amount of forced saving declines, consumers must partially replace the effective precautionary buffer that the forced saving provided by reacting more with their discretionary precautionary behavior.

These results supply a potential theoretical underpinning for the idea that the runup in consumer debt in the late 1980s was at least partly responsible for the severity of the decline in consumer spending, particularly on durable goods, in the recession and subsequent slow recovery in the early 1990s. However, the rise in debt-to-income ratios is not, in this interpretation, the driving force in the story; rather, both the increase in debt and an increased sensitivity of durables spending to unemployment expectations are emergent properties of the new stochastic equilibrium with deregulated credit markets.

It if worth emphasizing here how surprising this theoretical result is. The usual economic intuition is that relaxation of liquidity constraints should allow consumers to smooth consumption more. Here, a relaxation in liquidity constraints has exactly the opposite effect.

6.2 EMPIRICAL EVIDENCE

We turn, finally, to the question of whether there is any empirical evidence for the proposition that in the wake of financial deregulation durables spending has become more sensitive to unemployment expectations. We first perform the simplest possible test by examining whether the coefficient on the unemployment expectations variables has been significantly higher in the post-deregulation period than in the pre-deregulation period. The principal difficulty in performing this test is in deciding from when to date the deregulation. The initial stages of deregulation took place in the late 1970s during the Carter administration, and the policy reforms were largely complete by 1983. However, arguably the most important development (at least from the standpoint of its effect on the availability of mortgage credit) in the liberalized market was the rapid growth of the secondary market for mortgage debt fostered by

Table 11 INTERACTION TERM IN DURABLES REGRESSION[a]

Row	UE_t	UE_t^{Post85}	ΔUE_t	ΔUE_t^{Post85}	$\bar{R}^2$	D–W
1	−2.320		0.723		0.43	0.54
	−(6.19)***		(1.06)			
2	−2.032	−2.144	0.385		0.46	0.55
	−(6.39)***	−(2.60)***	(0.66)			
3	−2.341		0.486	1.322	0.43	0.54
	−(6.28)***		(0.71)	(0.85)		
4	−2.051	−2.076	0.283	0.626	0.46	0.56
	−(6.43)***	−(2.34)**	(0.42)	(0.38)		

[a]Dependent variable is the ratio of durables consumption to annuity labor income, 1963:2–1994:3. *, significant at 10% or better; **, at 5% or better; ***at 1% or better. *t*-Statistics are listed in parentheses below coefficient estimates. Standard errors were constructed using a serial-correlation-robust covariance matrix (allowing serial correlation at lags up to 18). UE_{t-1} is the unemployment expectations index, and UE_t^{Post85} is the index times a dummy variable equal to one from 1985:1 to the end of the sample period. The following were also included as independent variables but not reported: a constant term, the ratio of lagged annuity income to current annuity income (A_{t-1}/A_t), the prime rate ($Prime_t$), current income over annuity income (Y_t/A_t), and household net worth (NW_t/A_t). A constant term was also included but is not reported.

the Federal National Mortgage Association and similar government-sponsored enterprises. The associated rapid growth in mortgage debt appears to have begun around 1985. We therefore date the postliberalization period as beginning in 1985 (although our empirical results are not sensitive to the exact dates we choose).

Results are presented in Table 11. The interaction term on the level of unemployment expectations is highly statistically significant, and implies that the coefficient on unemployment expectations was roughly twice as large in the postliberalization period as in the earlier period.[38] However, the coefficient on the interaction term is insignificant for the variable measuring the change in unemployment expectations, once again reflecting our general empirical finding that the change in unemployment expectations is not nearly as reliably important as the level in influencing consumption choices.

7. *Conclusions*

The broad goal of this paper has been to document and then explain the relationships between household balance sheets and consumer pur-

38. We found similar results when we allowed all regression coefficients (not just the coefficients on the UE terms) to differ before and after 1985.

chase decisions. In our empirical work we found that unemployment expectations appear to have an influence on spending decisions beyond any information those expectations contain about future levels of income. We therefore develop a theoretical model of debt-financed durables purchases which has a serious role for labor-income uncertainty. This model implies that the location of the lower (S,s) trigger depends on the degree of labor-income uncertainty; when uncertainty increases, consumers postpone durables purchases until their balance-sheet condition improves. We find that this model does a much better job than the standard certainty-equivalent or fixed-band (S,s) models at explaining the cyclical dynamics of spending and balance sheets. However, the model does highlight a paradox: It is the lagged level of unemployment expectations, rather than the change in expectations (which the model emphasizes), that appears to be related to spending decisions.

This paper suggests a variety of important directions for future work. First, the calibration of the model was necessarily ad hoc. There appear to be surprisingly few data available about such important questions as how the risk of job loss changes over the business cycle, or how the job-finding hazard changes for those who are unemployed. Given the apparent empirical and theoretical importance of labor income uncertainty, this is an area where very useful work could be done. Second, the analysis of this paper treated unemployment expectations and the aggregate economic state as exogenous. Although in the wake of the rational-expectations revolution in macroeconomics it sounds staggering to say it, to our knowledge there has been virtually no recent research on how consumers' observable expectations are determined, either for the unemployment expectations variable we consider or for any of the other aggregate measures of consumer expectations. There are presumably many tests that could be performed to determine, for example, the rationality of those expectations. Fourth, the extreme short-term response of durables spending to uncertainty clearly raises the possibility of multiple equilibria in a general equilibrium version of this model. Although solving the full model in a general equilibrium setting is clearly well beyond current computational capacities, it is possible that simplified models which build in an extreme sensitivity of durables spending to uncertainty might be solvable. Finally, the model has many implications that are testable with microeconomic data. For example, a straightforward test would be to estimate a probit model of home purchase decisions and to test whether the purchase decision is affected by either local or aggregate unemployment expectations.

REFERENCES

Adda, J., and R. Cooper. (1997). Balladurette and juppette: A discrete analysis of scrapping subsidies. Boston: Department of Economics, Boston University. Manuscript.

Bar-Ilan, A., and A. S. Blinder. (1992). Consumer durables: Evidence on the optimality of usually doing nothing. *Journal of Money, Credit, and Banking* 24(2):258–272.

Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *Quarterly Journal of Economics* 98(1):85–106.

———. (1984). Permanent income, liquidity, and expenditure on automobiles: Evidence from panel data. *Quarterly Journal of Economics*, 99(3):587–614.

———. (1985). Adjustment costs, durables, and aggregate consumption. *Journal of Monetary Economics* 15(1):41–68.

———, G. Ferri, and P. Simon. (1997). Is there a flight to quality in consumer lending? Princeton, NJ: Princeton University. Manuscript.

———, M. Gertler, and S. Gilchrist. (1996). The financial accelerator and the flight to quality. *Review of Economics and Statistics* 78(1):1–15.

Bertola, G., and R. J. Caballero. (1990). Kinked adjustment costs and aggregate dynamics. In *NBER Macroeconomics Annual 1990*, O. J. Blanchard, and S. Fischer (eds.). Cambridge, MA: The MIT Press, pp. 237–288.

———, and ———. (1994). Irreversibility and aggregate investment. *Review of Economic Studies* 61(2):223–246.

Blanchard, O. J. (1993). What caused the last recession? Consumption and the recession of 1990–1991. *American Economic Review* 83(2):270–274.

Caballero, R. J. (1993). Durable goods: An explanation for their slow adjustment. *Journal of Political Economy* 101(2):351–384.

———, and E. M. R. A. Engel. (1994). Explaining investment dynamics in U.S. manufacturing: A generalized (*S,s*) approach. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 4887.

Campbell, J. Y. (1987). Does saving anticipate declining labor income? An alternative test of the permanent income hypothesis. *Econometrica* 55:1249–1273.

———, and A. S. Deaton. (1989). Why is consumption so smooth? *Review of Economic Studies* 56(3):357–374.

———, and N. G. Mankiw. (1989). Consumption, income, and interest rates: Reinterpreting the time-series evidence. In *NBER Macroeconomics Annual*, O. J. Blanchard and S. Fischer (eds.). Cambridge, MA: The MIT Press, pp. 185–216.

Caplin, A., and J. Leahy. (1997). Durable goods cycles. Boston: Boston University. Manuscript.

Carrington, W. J. (1993). Wage losses for displaced workers. *Journal of Human Resources* 28(3):435–462.

Carroll, C. D. (1992). The buffer-stock theory of saving: Some macroeconomic evidence. *Brookings Papers on Economic Activity* 1992(2):61–156.

———. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics* 112(1):1–56.

———, and W. E. Dunn. (1997). Data sources and solution methods for empirical and theoretical results in "Unemployment expectations, jumping (*S,s*) triggers, and household balance sheets." http://www.econ.jhu.edu/ccarroll.

———, J. C. Fuhrer, and D. W. Wilcox. (1994). Does consumer sentiment fore-

cast household spending? If so, why? *American Economic Review* 84(5):1397–1408.

———, and M. S. Kimball. (1996). On the concavity of the consumption function. *Econometrica*, 64(4):981–992.

———, and A. A. Samwick. (1997). The nature of precautionary wealth. *Journal of Monetary Economics* 40(1), August.

Deaton, A. S. (1991). Saving and liquidity constraints. *Econometrica* 59:1221–1248.

Dixit, A. K., and R. S. Pindyck. (1994). *Investment under Uncertainty.* Princeton, NJ: Princeton University Press.

Eberly, J. C. (1997). Flexible (*S,s*) bands, uncertainty, and aggregate consumer durables purchases. Philadelphia: University of Pennsylvania. Working Paper.

Fannie Mae Foundation. (1997). *Opening the Door to a Home of Your Own.* Washington.

Fratantoni, M. C. (1996). Housing wealth, precautionary saving and the equity premium. Baltimore: The Johns Hopkins University. Manuscript.

Garner, C. A. (1996). Can measures of the consumer debt burden reliably predict an economic slowdown? Federal Reserve Bank of Kansas City *Economic Review* 81(4):63–76.

Greenspan, A. (1992). Statement to the U.S. House Committee on Ways and Means, December 18, 1991. *Federal Reserve Bulletin* 78(2):122–124.

———, and D. Cohen. (1997). Motor vehicle stocks, scrappage, and sales. Washington: Board of Governors of the Federal Reserve System. Manuscript.

Grossman, S. J., and G. S. Laroque. (1990). Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods. *Econometrica* 58(1):25–51.

Hall, R. E. (1978). Stochastic implications of the life-cycle/permanent income hypothesis: theory and evidence. *Journal of Political Economy* 86(6):971–987.

———. (1993). Macro theory and the recession of 1990–1991. *American Economic Review* 83(2):275–279.

Kimball, M. S. (1990). Precautionary saving and the marginal propensity to consume. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 3403.

Ludvigson, S. (1996). Consumption and credit: A model of time-varying liquidity constraints. Princeton, NJ: Princeton University. Ph.D. Thesis.

Mankiw, N. G. (1982). Hall's consumption hypothesis and durable goods. *Journal of Monetary Economics* 10(3):417–425.

McCarthy, J. (1997) Does the household balance sheet affect expenditures response to monetary policy? New York: Federal Reserve Bank of New York. Manuscript.

Mishkin, F. S. (1976). Illiquidity, consumer durable expenditure, and monetary policy. *American Economic Review* 66(4):642–654.

———. (1977). What depressed the consumer? The household balance sheet and the 1973–75 recession. *Brookings Papers on Economic Activity* 1977(1):123–164.

———. (1978). Consumer sentiment and spending on durable goods. *Brookings Papers on Economic Activity* 1978(1):217–231.

Overland, J. (1997). Optimal saving with stochastic income and habit formation. Providence: Brown University. Manuscript.

Sichel, D. S. (1993). Business cycle asymmetry: A deeper look. *Economic Inquiry* 31(2):224–236.

———. (1994). Inventories and the three stages of the business cycle. *Journal of Business and Economic Statistics* 12(3):269–277.
Wilcox, D. W. (1992). The construction of U.S. consumption data: Some facts and their implications for empirical work. *American Economic Review* 82(4):922–941.

Comment

ANGUS DEATON
Princeton University

I am very much a fan of the line of work that Chris Carroll and his collaborators have been exploring in recent years. Precautionary saving is different from and much more interesting than we used to think. It can account for a range of interesting phenomena that are hard to explain within the standard certainty-equivalence version of life-cycle or permanent-income theory. In particular, precautionary behavior changes the role of assets, so that it makes sense to look to precautionary motives for a coherent account of the Greenspan hypothesis that the debt buildup in the 1980s contributed to the recession of the early 1990s. I think that the study has a good deal to it, and I am sympathetic to the general thrust of the paper. I like the emphasis on durable goods as well as on nondurables and assets, and I like and find plausible the proposition that financial deregulation has increased the vulnerability of the economy by making consumption more responsive to shocks. But this paper documents these general points in what seems to me a rather strange way, looking at some stylized facts while ignoring others, and constructing a model of housing that—while interesting in its own right—is of unproven relevance for the important hypothesis with which the paper begins; namely, that "deteriorated balance sheets" increase the propensity of the economy to fall into recession.

The paper begins with some regressions, based on various extensions of Campbell and Mankiw's (1989) consumption-growth equations. Nondurable consumption growth depends on the one-period-ahead anticipated rate of growth of income and on lagged unemployment expectations (more or less as in Carroll, Fuhrer, and Wilcox, 1994, although the variable used there was consumer sentiment). Various balance-sheet measures add little or nothing significant to these equations. Similar regressions are run for the ratio of consumption of durables to annuity income, for total home sales, and for the growth in total household liabilities. Temporarily putting these results to one side, Carroll and Dunn then develop an intertemporal choice model with two goods: nondurables

and housing services. Consumers can own or rent houses (at a premium); buyers are constrained to purchase houses whose value is three times their income, and must meet a down-payment constraint and pay transaction costs. Marginal utility is strictly convex, consumers experience spells of unemployment, and (the macroeconomic component) of income growth follows a serially correlated process that mimics the business cycle. The policy functions are solved out, and macro results obtained by aggregating simulations for 20,000 consumers. The data so produced are then subjected to the same battery of extended Campbell–Mankiw tests which began the paper, and the results compared. The model scores a few points, but by and large, the match is unimpressive. Indeed, several of the most important stylized facts are missed, including the correlations between consumption growth and both predictable income growth and lagged unemployment expectations.

Apart from the substantive results—to which I will return below—I have several concerns about this methodology. First, although the Campbell–Mankiw results provide a useful set of stylized facts that competing models should be able to explain, it is unclear why matching those results is an adequate substitute for fitting the model to the data. Carroll and Dunn's methodology picks a few correlations that have to be fitted and ignores all the others. For example, there is no attempt to check whether the time-series patterns of the simulations bear any resemblance to those in the data, although at least some of the tables (e.g., Table 9) suggest that they do not. Such a narrowly focused estimation strategy provides neither an adequate test of the model nor adequate recognition of what it can explain, even when it uniquely does so. A more thorough estimation and testing procedure might even provide positive evidence to offset the model's obvious deficiencies.

My second concern is a more general one and concerns how this sort of work should be reported. It is very hard to find out *exactly* what the authors did, even on a careful reading. The calculation of the policy function requires four days on a Unix workstation, and those of us who have performed similar (albeit much simpler) calculations know how hard it is to persuade oneself—let alone anyone else—that the calculations have been correctly performed. Matters are at their worst—as here—when there are no analytical results against which the calculations can be checked. Not only that, but there are no results here that guarantee the existence of policy functions or (conditional on policy functions) the existence of an invariant distribution or convergence to it. Models of precautionary saving are delicate, and we know from special cases of Carroll and Dunn's model that the existence of a limiting distribution depends on the values of the parameters. The complexity of the

calculations also precludes any serious attempt to choose the parameters so as to fit the data best. Hence, even if we stipulate the correctness of the calculations, we have no way of knowing whether the poor fit of the model comes from an unfortunate calibration, or rather is fundamental to the model's structure. The authors are really groping in the dark, and they expect a great deal of their audience when they ask us to accept that what they have here is what they say they have found. Of course, these problems are common to a great deal of applied work that roots itself in dynamic programming. But if the profession does not find a way of making these calculations credible and reproducible, their usefulness remains in serious doubt.

Nevertheless, I think the general approach is correct, and that precautionary motives can help explain the role of balance-sheet variables in affecting the rate of growth of consumption. Indeed, the sort of effects I have in mind have been previously emphasized in Carroll (1997), who showed that, in his buffer-stock model of precautionary saving, the conditional variance of consumption is inversely related to the level of assets. When the buffer stock is low, negative shocks to earnings and returns have larger consequences, the growth rate of consumption is expected to be more variable, and current consumption is restrained. More formally, the same result holds in the no-borrowing buffer-stock model developed in my own work (Deaton, 1991). In Deaton and Laroque (1992), we show that, with no borrowing allowed—in Carroll's model, there is voluntary abstention from borrowing, with much the same consequences—if (1) instantaneous marginal utility is convex in consumption, (2) the rate of interest is lower than the rate of time preference, and (3) earnings are i.i.d., then, conditional on information available at time $t - 1$, the variance of the value of the marginal utility of consumption at t is monotone nonincreasing in the level of assets. Provided that marginal utility is not too convex, the one-period-ahead variance of consumption will behave in the same way. In this case, assets are a precautionary buffer that ensures consumption *ex ante*. When precautionary stocks are low, consumers are vulnerable, for example to unemployment shocks, and an unanticipated negative shock will have a larger negative impact on consumption—and on the economy—than would have been the case had consumers been better cushioned with more assets. Consumption is rendered more sensitive to adverse shocks by a runup of debt, even one that was warranted on previous information. If regulation prevents consumers from assuming at least some of this debt, or makes them hold more assets than they wish, deregulation will increase the vulnerability of consumption.

Another version of much the same story can be tied to housing, al-

though through a different route from that developed in the paper. As has been emphasized by Modigliani in the context of life-cycle saving, precautionary motives may not be very important if wealth is being accumulated for other purposes. Modigliani argues that retirement balances can do double duty as precautionary balances, so that there is no separate role for precautionary saving; a corollary is that the replacement of life-cycle wealth with social security—which cannot be used as collateral—will make precautionary motives more important. Although I would argue that the importance of life-cycle saving is overstated in such arguments, the argument can be applied to housing. To the extent that deregulation has made it easier for people to buy houses, for example by lowering down-payment ratios, or raising the ratio of mortgage debt to income, there will be less saving up to buy a house, and less ability to go further into debt to deal with a bad shock. People will hold fewer assets (or more debt), and they will be made better off by not being forced to hold assets that they do not want, but their consumption will be more variable and more vulnerable to negative shocks.

Theoretical results of this sort make precautionary saving and deregulation plausible candidates for explaining the increased sensitivity of consumption to bad news. Of course, we are still a long way from having a precise and empirically supported account of such a mechanism. Carroll and Dunn's paper is a useful first step in this direction.

REFERENCES

Campbell, J., and N. G. Mankiw. (1989). Consumption, income, and interest rates: Reinterpreting the time-series evidence. In *NBER Macroeconomics Annual 1989*, Olivier J. Blanchard and Stanley Fischer (eds.). Cambridge, MA: The MIT Press, pp. 185–216.

Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics* 42:1–56.

Carroll, C. D., J. C. Fuhrer, and D. W. Wilcox (1994). Does consumer sentiment forecast household spending? *American Economic Review* 84:1397–1408.

Deaton, A. S. (1991). Saving and liquidity constraints. *Econometrica* 59:1121–1248.

———, and G. Laroque. (1992). On the behavior of commodity prices. *Review of Economic Studies* 59:1–23.

Comment

SIMON GILCHRIST
Boston University

This paper is an ambitious attempt to explain the notable weakness in consumer durables spending during the 1990–1991 recession. The pro-

posed explanation is that this weakness resulted from households' reluctance to take on more debt in the face of increased uncertainty regarding future employment prospects. The idea here is that the last recession was unusual in the sense that it was the first recession (and hence first substantial increase in unemployment risk) to have occurred after the financial deregulation of the late 1970s and early 1980s and the associated rise in debt during most of the 1980s. More generally, the paper provides a model of consumer behavior that links unemployment risk, durable and nondurable spending, and household balance sheets. The goal is to match some of the observed stylized facts regarding these variables, and to explore the implications of financial deregulation in the context of such a model.

Appropriately enough, the paper starts by documenting the stylized facts. The paper goes about documenting these facts using standard off-the-shelf regressions for nondurable and durable consumption expenditures. The paper also considers similar regressions for housing. The nondurables regression is of the Campbell–Mankiw variety, i.e., a regression of consumption growth on income growth and other variables that may matter in a world where the permanent-income hypothesis does not hold. The authors consider two types of variables as candidates by which to augment such regressions: consumer balance-sheet measures such as debt growth or debt service, and a consumer sentiment variable which measures consumers' expectations regarding the unemployment rate. The nondurable regressions are fairly persuasive in their findings: there appears to be no systematic relationship between nondurable consumption growth and balance-sheet variables, with the possible exception of debt growth, once one includes unemployment expectations in the regression. On the other hand, unemployment expectations have substantial predictive power for nondurables. The empirical analysis is thorough and persuasive with respect to this conclusion.

The durables regression relates consumer durable expenditures to annuity income growth and unemployment expectations. The baseline model is a variant of the Mankiw specification of consumer durable expenditures under the assumption of no significant transactions costs, adjustment costs, or other frictions. Again, this model finds a significant role for unemployment expectations. The paper also finds that lagged balance-sheet measures are positively correlated with durables expenditures, suggesting, unsurprisingly, an endogeneity regarding durables expenditures and debt.

While the paper is thorough in this analysis, it is not clear that these regressions are the most useful facts to document when matching the model to the data. The nondurables regressions are informative for comparison with other studies, but these are not necessarily the "moments"

one would want to match in a calibration exercise, particularly since there is no obvious link between the unemployment expectations measure collected by the University of Michigan and the uncertainty fed into the numerical model. From the perspective of durable-goods expenditures the focus on such regressions is more puzzling, given that the authors have in mind an (S,s) model of durables, which would not fit such regressions even if permanent-income behavior held. Perhaps more useful here would be to focus on simple correlations between variables rather than the more complicated regressions included in the paper.[1]

Although the empirical work takes up a substantial portion of the paper, the modeling section is where the innovations in the paper lie. The authors present a model of consumer behavior that incorporates nondurables and housing. The key elements to the model are lumpy housing and significant transaction costs to buying and selling houses. In addition, consumers face a substantial degree of idiosyncratic uncertainty through the process of unemployment. Capital markets are assumed to be imperfect in two important ways. First, there are no insurance markets which would allow consumers to diversify idiosyncratic risk, and second, there is a down-payment requirement of 20% on new-house purchases. The housing-market frictions imply (S,s) behavior for durable-goods expenditures. The combination of idiosyncratic risk and borrowing constraints through the down-payment requirements imply that consumers have a strong precautionary saving motive, and are willing to postpone their purchases of the durable good in order to save cash when unemployment risk increases. Since recessions are periods when consumers face high unemployment risk, a substantial fraction of consumers respond to the recession by delaying purchases—resulting in a large shift in the lower trigger of the (S,s) band for durables. Given this structure, calibrated to match certain features of U.S. data, the model is solved, simulated, and aggregated to produce statistics from which to run regressions and match results to aggregate data.

The intuition that consumers may delay purchases of durables in response to increased unemployment risk during the recession is very appealing, as is the notion of (S,s) bands responding because of the increased value of liquid assets at the onset of recessions. Unfortunately,

1. For example, the (S,s) literature on consumer durables often focuses on the degree of "excess smoothness" of durables expenditures relative to a benchmark frictionless model. Such smoothness is summarized by the MA coefficient on a univariate durables expenditures equation. Linking the degree of smoothness to the severity of financial constraints in the model would seem worthwhile.

the sheer size of the model and the complexity of calibration make many statements beyond that difficult. In addition, from a calibration point of view, it is very difficult to judge whether or not the model is a success.[2] It replicates some features of the data, most notably that unemployment expectations help forecast consumption growth (although the data suggest that the level should matter, whereas the model finds that only the growth rate should matter), but not others. Unfortunately, the authors provide very little information regarding how well the model matches basic correlations and autocorrelations of the key model variables. More perplexingly, the model does not seem to be capable of replicating the Campbell–Mankiw result that forecastable income growth predicts nondurables consumption growth.

After presenting model simulations, the paper focuses on an exercise intended to illuminate the effects of financial deregulation. The experiment considered is an exogenous reduction in the down-payment requirement from 20% to 10%. In the new steady state that results from this reduction, consumer spending on nondurables and durables is found to be more sensitive to unemployment risk. The intuition behind this result is that consumers have higher levels of debt and hence greater fixed payments to consider when deciding if and when to liquidate financial assets to buy a new house.[3] It is this increased sensitivity that the authors use as an explanation for the weakness of the durable goods expenditures during the 1990 recession. Unfortunately, as the authors mention, the degree of increased sensitivity seems very difficult to pin down and varies across parameter values, making any statement beyond "it can occur" somewhat difficult to make with a high degree of confidence.

Besides the difficulty in obtaining precise answers regarding the quantitative effects of deregulation as it is currently modeled, it seems very difficult to answer the question qualitatively without knowing more about how such deregulation affects both the supply and demand for consumer goods and the supply and demand for debt. In particular, one

2. Despite the model's complexity, it omits certain key features such as risk associated with asset values and variations in interest rates. The former is especially important if one can borrow against the collateral value of the house. To the extent that changes in interest rates ease debt burdens, the latter provides a significant link between monetary policy and consumer spending.
3. If the basic driving force behind the increased precautionary savings is the fact that consumers have a substantial part of their current income accounted for through fixed debt payments, one wonders whether there isn't a simpler model that would capture this effect and still provide the link between increased unemployment expectations and consumer spending.

would expect financial deregulation to have a large impact on asset values given constraints on supply in the short run. To the extent that consumers can borrow against the collateral value of their existing homes, such a rise in asset values could fuel a prolonged expansion in consumer credit. Under this scenario, the extended runup in debt and subsequent collapse of durables in the 1990–1991 recession may owe more to the dynamics of the transition to the new steady state than to any dynamics that may occur from an economy starting out in the new steady state.

More generally, the paper does not stop to consider exactly what type of financial deregulation has occurred throughout the 1980s, and what are its broader consequences and implications for consumer spending. The fact that debt levels have risen is well documented, as is the fact that debt service burdens rose steadily. The paper's view of financial deregulation is primarily a relaxation of down payment restrictions owing to innovations in the mortgage market. To investigate how realistic this is, I obtained data on loan-to-value ratios for houses, new cars, and used cars.[4] These are plotted in Figure 1. As one can see, there is some tendency for car loan-to-value ratios to rise during the 1980s but little tendency for housing ratios to do so. Thus the notion that the economy has settled at a new, higher level in terms of mortgage debt per household seems far fetched, as does the notion that the majority of house buyers face down-payment difficulties (this seems especially true in view of current down-payment requirements, which are only 5%). If loan-to-value ratios did not rise on average, but financial deregulation occurred in the housing market, the effect of such deregulation may have come from new consumers previously shut out of the market. If so, this suggests a completely different type of experiment should be run with the model.

While one can quarrel with the specifics of the financial deregulation exercise, this paper does makes some steps towards formalizing the link between consumer durables and the severity of financial frictions through the unwillingness of consumers to take on more debt in periods of high income uncertainty. In this model, such a link comes through a down-payment requirement. In today's world of readily available credit and 5% down payments, however, it would seem that actual credit constraints are unlikely to occur with great frequency for most households. Even if actual credit constraints do not occur, such households may still face credit frictions through high premiums on borrowing rates. To examine how high such premiums might be, I obtained data on

4. The source for these data is various issues of the *Federal Reserve Bulletin*.

Figure 1 LOAN-TO-VALUE RATIOS

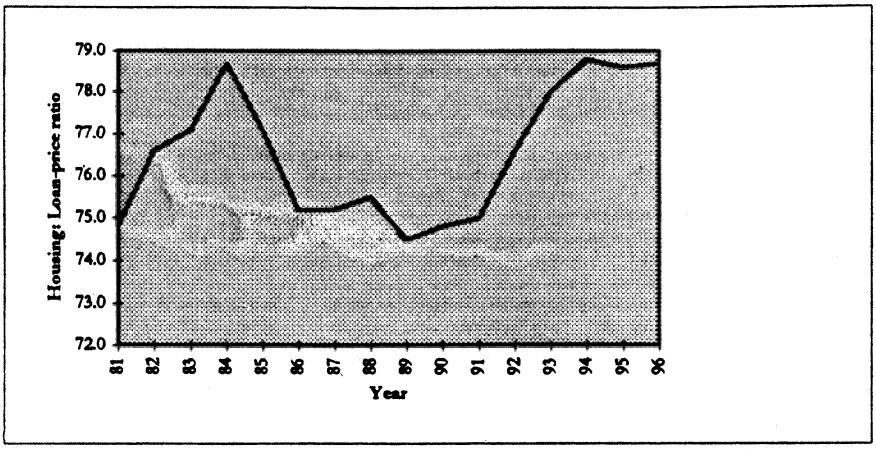

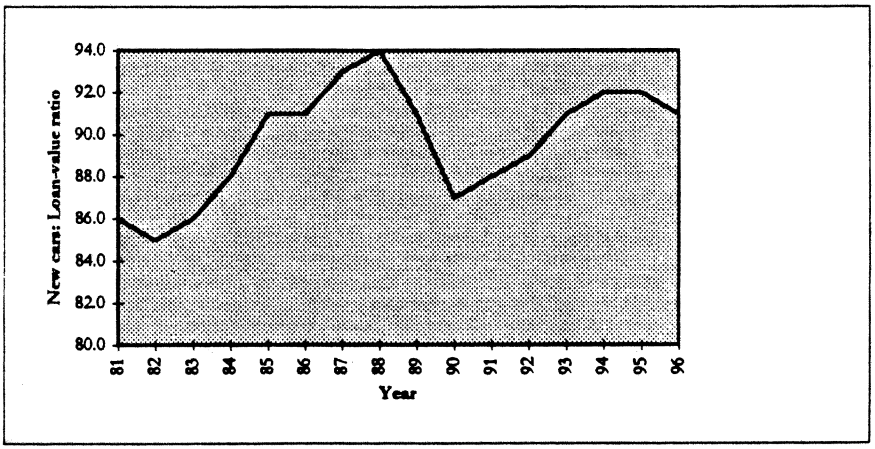

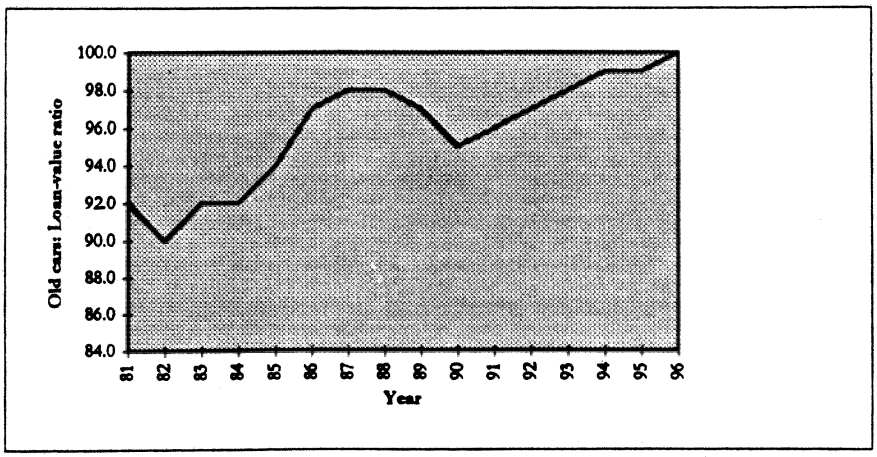

1996 values are for the month of June only.

Figure 2 CONSUMER RATE SPREADS

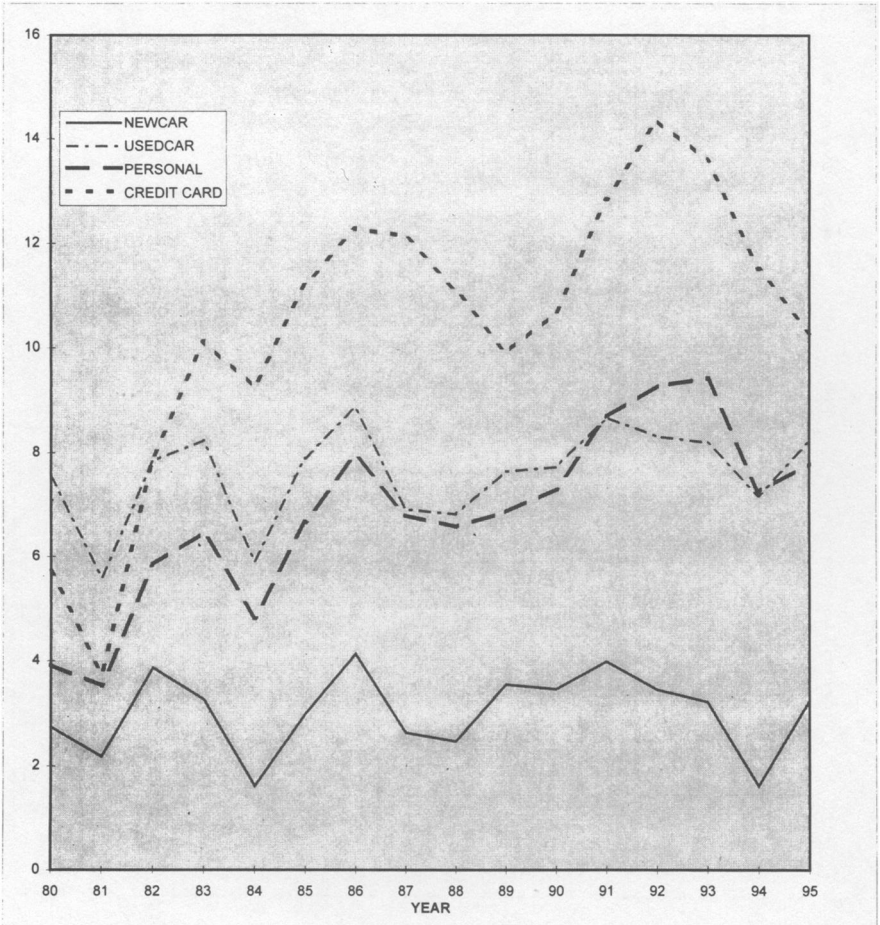

interest-rate spreads for (ordered by degree of collateral) new cars, used cars, 24-month personal loans, and credit-card debt.[5] These data are plotted on an annual basis in Figure 2. The plot reveals that the absolute levels of these rate spreads are high. For example, the 7.5% average spread of a used-car loan represents a 30% premium on the user cost of

5. The data were again obtained from various issues of the *Federal Reserve Bulletin*. The interest rate spreads were calculated by taking the difference betweent the actual interest rate and an appropriately matched safe T-bill or government bond rate. To decide on the relevant maturity for new and used car loans I used the average maturity numbers reported in the *Bulletin*.

capital for old cars.[6] For credit-card loans, the premiums over safe rates average 14%. In addition to being high on average, consumer rate spreads clearly vary inversely with the degree of collateral (new cars have the lowest spread, credit-card debt the highest) and tend to increase during business-cycle downturns. In addition, rate spreads have steadily risen in most categories throughout the 1980s, as indebtedness has increased.

These facts again suggest a slightly different view of credit liberalization, namely, that while more consumers obtain access to credit, they are paying a large premium to do so. More generally, such facts suggest that premia on external funds may be an important component of consumer decision making. Obtaining a better understanding of how important such premia are for the dynamics of consumer durables seems like a logical next step in what appears to be a rapidly evolving literature on consumer behavior.

Discussion

Chris Carroll began by commenting on a few points made by the discussants. In defense of the neglect of housing-price risk in their paper, Carroll cited work by Mark Frantantoni on the effect of housing-price risk on agents' decisions to hold risky assets. Frantantoni's paper found that one can separate the financial risks of homeownership into two parts: the risk associated with variations in the house price and the risk created by the fixed obligation of monthly mortgage payments, which reduces the homeowner's ability to smooth consumption and effectively increases risk aversion. Fratantoni found the latter effect to be the more important, justifying Carroll and Dunn's emphasis on that channel.

In response to comments that their paper focused too narrowly on housing purchases and financing, Carroll argued that housing-related assets and liabilities dominate consumer balance sheets, the analysis of which had been their original objective. Janice Eberly responded that, while she found the effects of quasifixed housing expenditures on consumer spending and saving behavior very interesting, she was not convinced that reliance solely on housing data was the best way to calibrate these effects. She noted that households have a variety of mechanisms

6. While it is difficult to sort out how much of the spread over the safe rate represents compensation for average default risk, the potential distortions created by such spreads appear to be large. For example, even on comparing the average spread for used-car loans with that for new-car loans (4.6%), one still obtains a 20% premium in the user cost of used vs. new cars.

for managing their home payments, such as varying the size of the house purchased or the structure of the mortgage terms. She suggested that it would be simpler to look directly at households' total fixed payments, which could be modeled as costly to adjust. Carroll noted that they had taken this approach to some extent in the first part of their paper in their analysis of the debt service burden, but they had found no evidence that the level of debt service per se was an important determinant of consumption.

Greg Mankiw noted that although he liked buffer-stock models, work by John Shea has led him to have some nagging doubts. Shea's research explores the prediction that households facing liquidity constraints should respond asymmetrically to expected changes in income. Shea found the asymmetry to be the opposite of that predicted by the theory: Specifically, he found that consumption does not respond to expected increases in income but does respond to expected income declines. Carroll contended that Shea's results are generally not very robust, a view that was seconded by Angus Deaton.

Anil Kashyap observed that there do exist companies that are willing to insure consumer mortgage payments against the risk of unemployment. He conjectured that two factors may lead to a lack of demand for this insurance—it may not be priced actuarially fairly, or consumers may already mitigate the unemployment risk by other means, such as precautionary saving or choosing a smaller house. Kashyap suggested that the paper needs to explain why people do not generally take this type of insurance. Carroll offered the possibility that this type of insurance might be relatively new and hence unfamiliar to consumers. In any case, he argued, in practice it seems difficult to discount the effects of unemployment risk, as the data suggest that unemployment expectations are very important for determining home sales and purchases of durable goods. Carroll conceded that their model does not fully explain the lack of risk sharing but emphasized that its ability to rationalize a role for unemployment expectations in the house purchase decision, independent of the expected level of income, is a desirable feature.

Robert Hall again raised the issue of house price risk; he wondered why this risk was particularly pertinent to current consumption decisions, given that the flow of housing services is unchanged, and holding constant current and expected future income. Deaton replied that the issue is that the equity will not be there if needed (i.e., if there is a bad income draw); indeed, one cannot sell the house and get out of the leveraged position at all if the home's value declines by enough.

Julio Rotemberg noted that legal changes have made it easier to seek refuge in bankruptcy now than in the past, and that perhaps this is a

reason people are taking on more debt. He suggested that more lenient bankruptcy laws could be thought of as increased insurance against poor income draws; and if so, these changes arguably would increase rather than decrease the empirical relevance of the permanent-income model. Carroll expressed skepticism of the view that people think of bankruptcy as a contingency plan, and so he did not agree that the legal changes were likely to be important factors.