AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Liquidity Constraints of the Middle Class*

Jeffrey R. Campbell[†] and Zvi Hercowitz[‡]

November 2009

Abstract

The consumption of households with liquid financial assets responds much more to transitory income shocks than the permanent-income hypothesis predicts. That is, middle class households with assets act as if they face liquidity constraints. This paper addresses this puzzling observation with a model of impatient households that face a large recurring expenditure. In spite of impatience, they save as this expenditure draws near. We call such saving made in preparation for a foreseeable event at a given future date “term saving”. Term saving reverses the role of assets in the presence of liquidity constraints. Typically, assets grow following past lucky events; thus assets imply an *abundance* of liquidity. Here, assets indicate an impending need for funds and a *shortage* of liquidity. The borrowing constraint will bind at the time of the expenditure, and assets will then be zero. This separates planning up to that time from the rest of the household’s lifetime and thereby shortens its effective horizon. Intertemporal substitution over a limited period generates a strong consumption response to temporary income changes. As the expenditure approaches, the effective horizon shortens further and the household accumulates assets. Hence, households with more assets have *larger* consumption responses. We compare a calibrated version of the model with observations from the 2001 U.S. tax rebate and with evidence on excess smoothness and persistence of aggregate consumption.

*We thank R. Andrew Butters, Ross Doppelt, and Ryan Peters for their excellent research assistance and Sumit Agarwal, Gadi Barlevy, and Mariacristina DeNardi for their thoughtful comments. The views expressed herein are those of the authors. They do not necessarily reflect the views of the Federal Reserve Bank of Chicago, the Federal Reserve System, or its Board of Governors.

[†]Federal Reserve Bank of Chicago, U.S.A.

[‡]Tel-Aviv University, Israel

JEL Code: E21

Keywords: Fiscal Policy, Tax Rebates, Marginal Propensity to Consume, Term Saving

1 Introduction

Middle-class households usually hold liquid assets. Since liquid wealth can be converted into current consumption, it seems implausible that such households could face liquidity constraints. Nevertheless, existing evidence from consumption dynamics following tax rebates suggests otherwise: Households with financial assets spend no less, and perhaps more, out of these temporary income changes than do poorer and more obviously constrained households. Thus, it appears that substantial liquidity constraints apply also to the middle class.

In this paper, we focus on the marginal propensity to consume, because of its link to the pervasiveness of liquidity constraints. We model middle-class households as home owners who hold assets. The model displays high marginal propensities to consume out of temporary income—even for households with liquid wealth. The key elements of our model are a motivation to save, impatience, and a limitation on debt. The motivation to save is a foreseeable life-cycle expenditure with exogenous timing and endogenous size. For tractability, we place this into a standard infinitely-lived household representation of a dynastic life-cycle model. Impatience implies that the household’s subjective discount rate exceeds the interest rate. In [Campbell and Hercowitz \(2009\)](#), the lower interest rate arises from the savings of more patient (and thereby wealthier) households. Here, we take the interest rate as given. The limit on debt mimics typical household debt contracts in the U.S., which require collateral in the form of a house or car with value exceeding the debt.

Impatience and the periodic expenditure pull the household’s assets in opposite directions, as impatience and risk do in [Carroll’s \(2001\)](#) model of precautionary saving. With impatience alone, the household would always hold the minimum required equity in its durable goods stock and have no financial assets. The certain periodic expenditure motivates this household to save. Such “term saving” finances a deterministic expenditure, while risk motivates precautionary saving. The two concepts of saving complement each other. In their pure forms, term saving is motivated by a foreseeable event at a given future date, while precautionary saving is generated by a constant risk. Hence, precautionary saving leads to a stationary buffer stock of assets; while term saving generates asset cycles. Below, we review evidence from the 2001 Survey of Consumer Finances which indicates that term saving for foreseeable major life-cycle events, such as college education and health care in old age, is at least as prevalent as precautionary saving among the middle class.

In the model’s conflict between impatience and saving, impatience wins when the periodic expenditure is far ahead—and thus the borrowing constraint binds. However, consumption smoothing eventually motivates the household to start saving ahead of the expenditure so

that the borrowing constraint ceases to bind. At the time of the expenditure, the household dissaves and the borrowing constraint binds again.

[Zeldes \(1984, 1989\)](#) distinguished between a *currently* binding liquidity constraint—the usual notion—and a *globally* binding constraint. The latter includes the possibility of the constraint binding in the future. As he noted, the expectation that a borrowing constraint will bind in the future effectively shortens the horizon over which a currently unconstrained household optimizes. This in turn generates a large *MPC* out of transitory income.

In the present model, all households are globally liquidity constrained, including those with assets, because the borrowing constraint necessarily binds at the time of the periodic expenditure. A household that is close to the periodic expenditure faces a shorter effective planning horizon. In the model’s calibrated version, households begin to accumulate assets well ahead of the next expenditure, and thus their Euler equations hold with equality most of the time. Furthermore, this asset accumulation implies that households with shorter effective planning horizons have more assets. So, the model predicts that for households with assets, the observed *MPC* increases with wealth.

The remainder of this paper proceeds as follows. In the next section, we review existing evidence on consumption behavior that points to liquidity constraints. This includes observations of the marginal propensity to consume out of tax rebates as well as the excess smoothness and stickiness of aggregate consumption growth (as defined by [Campbell and Deaton \(1989\)](#) and [Carroll, Slacalek, and Sommer \(2009\)](#)). Section 3 derives our key result analytically for a simple version of the model without durable consumption goods and a fixed debt ceiling. We then add durable goods and collateralized debt to the framework as preparation for the quantitative analysis. Section 4 considers the quantitative implications of a calibrated version of the model with regard to the facts reviewed in Section 2. In particular, we show that the liquidity constraints on the middle class can produce high *MPC*s for households with wealth and contribute to excess smoothness and stickiness of aggregate consumption. Section 5 offers concluding remarks.

2 Consumption Evidence

This section reviews existing evidence on consumption dynamics that motivates our exploration of middle-class liquidity constraints. Cross-sectional observations on households’ consumption decisions in the wake of tax rebates and changes indicate high marginal propensities to consume, especially for households with assets. Many of these households respond to

tax changes *as if* they were facing binding liquidity constraints. We also review some of the analyses of aggregate consumption that suggest widespread presence of liquidity constraints.

2.1 Cross-Sectional Evidence

Shapiro and Slemrod (2003) and Johnson, Parker, and Souleles (2006) used household surveys to collect evidence on the marginal propensity to consume out of the tax rebate in 2001 that began President Bush’s ten-year tax cut.¹ From July to October, the treasury mailed rebates of taxes paid under the previously higher rates to most taxpayers. Shapiro and Slemrod attached questions to the University of Michigan’s monthly Survey of Consumers, soliciting respondents’ anticipated use of these funds as well as their expectations about the evolution of future government spending and taxes. Johnson, Parker, and Souleles added a question in the Consumer Expenditure Survey to learn about rebate recipients’ marginal propensities to consume.

Both investigations found that consumers spent substantial fractions of their rebates. Johnson, Parker, and Souleles found that nondurable consumption increased by about two thirds of the rebate during a six-month period (close to 40 percent during the three-month period of the rebate, and the remaining in the following three months). Shapiro and Slemrod found that 22 percent of respondents anticipated spending most of the rebate, while the rest planned either to reduce their debts or increase their savings. Using plausible distributions of the marginal propensities to consume across those who would “mostly spend” and “mostly save”, Shapiro and Slemrod calculated an average marginal propensity to consume of about one third.

The original legislation specified that the tax cuts would sunset in 2011, but Congress could have either made them permanent or revoked them entirely before then. Shapiro and Slemrod’s respondents indicated that they expected the tax cuts to be temporary. When asked whether they expect smaller, the same size, or larger tax cuts in the future, 37 percent responded “smaller,” 47 percent replied “the same,” and only 16 percent responded “higher.” Regarding future government spending, 26 percent responded that they expect higher spending, 55 percent expect no change, and only 19 percent expected smaller spending. Hence, 81 percent expected the same or higher government expenditures. If households recognize the link between spending and taxes—or, in other words, if they have “Ricardian beliefs”—they should also expect a short-lived tax cut to engender a future tax increase.

¹This rebate authorized by the Economic Growth and Tax Relief Reconciliation Act of 2001.

Stock Ownership Class	Percentage of of Sample	Percentage Spending Most of Rebate
None	42.8	19.5
\$1 – \$15,000	9.1	13.1
\$15,001 – \$50,000	9.9	18.1
\$50,001 – \$100,000	6.8	26.7
\$100,001 – \$250,000	6.2	33.6
More than \$250,000	5.1	22.9
Refused/Don't Know	20.1	25.3

Table 1: Rebate Spending Percentages from [Shapiro and Slemrod \(2003\)](#)

Source: Table 2 of [Shapiro and Slemrod \(2003\)](#)

However, [Shapiro and Slemrod](#) found little evidence of a Ricardian link between expectations about future government spending and consumption responses.

When asset ownership indicates the absence of liquidity constraints, we expect the marginal propensity to consume to decline with the level of assets. The surprising finding from both investigations is quite different. [Shapiro and Slemrod](#) sort their sample by stock ownership, the only data on wealth in their survey. They divide their sample into six groups, and Table 1 reports their results. Remarkably, the spending fraction *increases* with stock ownership, with exceptions for the highest bracket and that with zero-assets. They find the same pattern using income levels. [Shapiro and Slemrod](#) report in their article's original working paper that this pattern also arises in regressions with dummy variables for the different stock ownership brackets, while age and other control variables are included. However, the relationship is statistically indistinguishable from a flat line.²

[Johnson, Parker, and Souleles](#) sort their sample into three groups by liquid assets. Households in their low-asset group spent much more than those in the middle-asset group, but those with the most liquid assets also spent more than those in the middle.³ These authors find the same pattern after sorting by income, which might proxy for other financial assets.

Looking beyond the 2001 tax rebates, there is additional evidence suggesting that marginal

²See their Tables 10 through 13 of NBER Working Paper 8672.

³See their Table 5. Note that the latter difference between the middle and high asset groups is not statistically significant.

propensities to consume out of transitory income are positively correlated with financial assets, rather than negatively. [Parker \(1999\)](#) examines households in the Consumer Expenditure Survey after their earnings pass the Social Security tax wage cap. This produces a forecastable and temporary increase in after-tax earnings for high-income families. He compares two groups of households, one with enough assets for at least 6 months of nondurable consumption and another with insufficient assets for one month of nondurable consumption. His estimates of the consumption elasticity with respect to present earnings for the two groups are 0.8 and 0.5. Only the estimate for the high-asset group is statistically significant.

[Souleles \(1999\)](#) documents the consumption responses of Consumer Expenditure Survey respondents to their annual tax refunds. This cash inflow is both predictable and transitory. Hence, under the permanent-income hypothesis, there should be no response. When he partitions the sample by the ratio of wealth to earnings, he finds that the least wealthy have a small and statistically insignificant response of total consumption, while the wealthiest spend between twenty-one and twenty-six cents for every dollar of the rebate. This spending was apparently concentrated in durable consumption, because the results reverse when examining food consumption and the “strictly nondurable” consumption defined by [Lusardi \(1996\)](#). Nevertheless, [Souleles \(1999\)](#) reports finding substantial responses of nondurable consumption among the wealthy to tax rebates after one quarter.

This body of evidence indicates that many households with assets act as if they are liquidity constrained. We base our explanation for these observations on households’ rational processing of available information. An alternative is a behavioral explanation of the evidence based on imperfect incorporation of information into current decisions. This paper is not aimed at discriminating between these two possibilities. Rather, it only provides a rational-expectations benchmark model for interpreting these data.

2.2 Time-Series Evidence

The cross-section evidence on consumption reactions to income changes reviewed above has a direct time-series counterpart in [Flavin \(1985\)](#), who documents “excess sensitivity” of consumption growth to current income growth. The basic difference between these two approaches is that tax cuts are clearly exogenous to the individual household, while in aggregate time-series current income changes have to be instrumented. Nevertheless, the cross-sectional evidence is consistent with [Flavin’s](#) findings.

[Deaton \(1987\)](#) provides a complementary perspective on aggregate consumption’s reaction to income changes that also quantifies failure of the permanent-income hypothesis. He

notes that if income growth is positively serially correlated, permanent income growth is more volatile than income growth. In this case, the permanent-income hypothesis predicts that consumption growth should be more volatile than income growth. The empirical results in [Deaton \(1987\)](#) and [Campbell and Deaton \(1989\)](#) indicate the opposite. Although income growth is indeed positively serially correlated, its standard deviation substantially exceeds that of consumption growth. In this sense, there is “excess smoothness” in consumption. The positive serial correlation of consumption growth, which [Carroll, Slacalek, and Sommer \(2009\)](#) label “consumption stickiness,” is related to excess smoothness: Both features follow from sluggish consumption movements that contradict the white noise prediction of the permanent-income hypothesis. [Carroll, Slacalek, and Sommer \(2009\)](#) estimate an autocorrelation coefficient of 0.83 for quarterly consumption growth in the United States.⁴ Its annual counterpart equals 0.62.⁵

Habit formation provides the most straightforward qualitative explanation for excess smoothness and consumption stickiness, but [Dynan \(2000\)](#) finds little evidence of habits in the PSID’s observations of food consumption. Of course, habits contribute nothing to reconciling the high *MPCs* reviewed above with the presence of wealth. In contrast, the periodically binding liquidity constraints we explore generate (relative to the predictions of the permanent-income hypothesis) both excess sensitivity to a transitory income change and excess smoothness and persistence following a permanent income change.

3 The Model

Our model of middle-class consumption and savings decisions adds a motivation to save to the impatient, borrowing-constrained household in [Campbell and Hercowitz \(2009\)](#). For this, we give the household utility from a special expenditure with predetermined timing but endogenous size. This specification keeps the resulting model simple, and it has the added benefit of realism. In 2001, respondents to the Survey of Consumer Finances were asked:

In the next five to ten years, are there any foreseeable major expenses that you and your family expect to have to pay for yourselves, such as educational expenses, purchase of a new home, health care costs, support for other family members, or anything else?

⁴See the left-most column of their Table 1.

⁵For this calculation, we presumed that quarterly consumption growth follows a first-order autoregression and calculated the first autocorrelation of its annual average.

Among middle-class households with working-age heads—a sample we define more precisely below in Section 4.1—58 percent answered affirmatively. Of those, 61 percent were saving for the expense at the time of the interview. Put together, about 35 percent of middle-class households with working-age heads were saving in anticipation of a foreseeable major expense. This percentage is similar to the fraction of households reporting that they sometimes save for precautionary reasons, 33 percent.⁶

To keep the model as simple as possible, we interpret “foreseeable” as “certain on a given date.” The model household derives utility from a special periodic expenditure and from regular consumption. We intend the periodic expenditure to represent major expenses as in the question above.⁷ The household lives forever and is impatient relative to the market rate of interest. In spite of the impatience, the household saves in anticipation of the periodic expenditure. Recall from the introduction that a household is globally liquidity constrained if its borrowing constraint either binds in the present or is expected to bind in the future. Term saving induced by the periodic expenditure generates a meaningful difference between current and global liquidity constraints: As we show later, the borrowing constraint should bind at least at the time of the expenditure. Hence, to different degrees, the household will always be liquidity constrained in the global sense.

Since the vast majority of household debt is collateralized by durable goods, such as homes and cars, explicitly accounting for collateral increases the empirical relevance of any quantitative investigation of liquidity constraints. Nevertheless, the main features of our model can be shown with nondurable consumption only. Accordingly, we first derive the key results qualitatively from that basic version of the model, and then conclude this section by extending it for the quantitative investigation.

⁶For this calculation, we used responses to the question “Now I’d like to ask you a few questions about your family’s savings. People have different reasons for saving, even though they may not be saving all the time. What are your family’s most important reasons for saving?” Each respondent can give up to six answers. We assigned precautionary motives to a household if any of these answers was: “Reserves in case of unemployment,” “In case of illness; medical/dental expenses,” “Emergencies; ‘rainy days’; other unexpected needs; for ‘security’ and independence,” or “Liquidity; to have cash available/on hand.” The fraction of households that reported precautionary motives for saving and also said they were saving for a foreseeable major expenditure was 11 percent.

⁷The Appendix presents an explicit overlapping-generations life-cycle model with both the periodic life-cycle expenditure and tax-advantaged retirement accounts. In Subsection 3.6 below we review that model and comment on its equivalence to the present one.

3.1 The Basic Model

The model proceeds in discrete time, and we denote a point in time as a “year”. A single infinitely lived household values two goods, standard consumption and periodic consumption. We denote the quantities of these consumed in year t with C_t and M_t . The utility function is

$$\sum_{t=0}^{\infty} \beta^t \{\ln C_t + \mu_t \ln M_t\} \quad (1)$$

with $0 < \beta < 1$. Here, the indicator μ_t follows a cycle with $\mu_t = \mu > 0$ every τ years and $\mu_t = 0$ at other times. This specification generates a periodic expenditure with exogenous timing and endogenous size. [De Nardi, French, and Jones \(2009\)](#) use a similar preference specification to study endogenous medical expenditure choices in old age.

The household is endowed with one unit of labor which it supplies inelastically at the wage rate W_t . Denote lump-sum taxes with T_t and net financial assets at the end of the previous year with A_t . The household’s budget constraint is

$$C_t = W_t - T_t + RA_t - A_{t+1} - M_t, \quad (2)$$

where R is the gross interest rate, assumed to be constant. We assume that $\beta R < 1$, so the household is impatient.⁸

The household’s choices of all goods must satisfy nonnegativity constraints. Furthermore, the household faces the standard borrowing constraint

$$A_{t+1} \geq 0. \quad (3)$$

Given A_0 , the household chooses sequences of C_t , M_t and A_{t+1} to maximize its utility subject to the sequence of nonnegativity and budget constraints.

Denote the Lagrange multipliers on the year t budget and borrowing constraints with Ψ_t and Γ_t . The first-order conditions for optimization are

$$\Psi_t = \frac{1}{C_t}, \quad (4)$$

$$\Gamma_t = \Psi_t - \beta R \Psi_{t+1}, \quad (5)$$

$$\Psi_t M_t = \mu_t. \quad (6)$$

⁸See [Campbell and Hercowitz \(2009\)](#) for a general equilibrium environment in which such a low interest rate arises endogenously from trade with a more patient household.

Without borrowing constraints, Ψ_t equals the marginal current utility of lifetime resources. Here, it represents the marginal value of current resources. The multiplier Γ_t equals the marginal value of relaxing the borrowing constraint, which is the deviation from the standard Euler equation; Γ_t is zero when the borrowing constraint is slack. Because Ψ_t is always positive, the periodic expenditure M_t is positive when $\mu_t > 0$ and zero otherwise. We elaborate on the life-cycle interpretation of M_t below in Section 3.6.

3.2 The Nonstochastic Cycle

Because of the periodic changes in preferences, the household's problem has no steady state, even if wages and taxes remain unchanged. Nevertheless, there does exist a nonstochastic cycle when W_t and T_t are constant. This cycle is the analogue of a steady state in our model, so we begin with the cycle's characterization, and focus in particular on term saving, i.e., the level of assets along the cycle. For this, we denote ordinary consumption and assets κ years after the most recent periodic expenditure in a nonstochastic cycle with C^κ and A^κ .⁹

From (4) and (5), the necessary conditions which a nonstochastic cycle must satisfy are

$$\frac{1}{C^\kappa} \geq \frac{\beta R}{C^{\kappa+1}} \text{ for } \kappa = 1, 2, \dots, \tau - 1, \quad (7)$$

$$\frac{1}{C^\tau} \geq \frac{\beta R}{C^1}. \quad (8)$$

The corresponding budget constraints are

$$\begin{aligned} C^\kappa + A^{\kappa+1} &= W - T + RA^\kappa \text{ for } \kappa = 1, 2, \dots, \tau - 1, \\ (1 + \mu)C^\tau + A^1 &= W - T + RA^\tau. \end{aligned}$$

We replaced here the periodic expenditure with its optimal level from (4) and (6), μC^τ .

To solve these conditions, it is helpful to begin with the case of $\mu = 0$, which corresponds to the standard optimization under impatience. The only path for A^κ and C^κ satisfying these conditions is the standard steady state for impatient agents in which the borrowing constraint always binds and the household consumes all labor earnings. That is, $A^\kappa = 0$ and $C^\kappa = W - T$ in all periods, and hence (7) and (8) hold with strict inequalities.

⁹Our model has a deterministic asset cycle in common with the models of [Baumol \(1952\)](#) and [Tobin \(1956\)](#). This and those models, however, differ in key respects. There, the length of the cycle is the key endogenous variable, while here it is exogenous. We focus is on the link between the asset cycle and liquidity constraints, while those models focused on the link between assets and money demand.

Raising μ above zero generates a positive M every τ years. However, if μ is less than $\hat{\mu}^\tau \equiv (\beta R)^{-1} - 1$, then the borrowing constraint still binds at all times. That is, for $\kappa = 1, \dots, \tau - 1$, $C^\kappa = W - T$, and $C^\tau = (W - T) / (1 + \mu)$. Thus, conditions (7) and (8) still hold with strict inequalities. Mechanically, this follows from the fact that the reduction of C^τ as μ goes up is not enough to bring

$$\frac{1}{C^{\tau-1}} \geq \frac{\beta R}{C^\tau} \quad (9)$$

to an equality while $C^{\tau-1}$ still equals $W - T$. Intuitively, the anticipated reduction in consumption is too small to induce the household to save in year $\tau - 1$ towards the expenditure in year τ , so the household finances the expenditure only by reducing C^τ .

Now, suppose that $\mu > \hat{\mu}^\tau$ and define $\hat{\mu}^{\tau-1} \equiv (\beta R)^{-2}(R + 1) - 1 > \hat{\mu}^\tau$. When $\mu = \hat{\mu}^{\tau-1}$, then the saving for the periodic expenditure reduces $C^{\tau-1}$ to $\beta R(W - T)$, so that:

$$\frac{1}{C^{\tau-2}} = \frac{\beta R}{C^{\tau-1}}. \quad (10)$$

That is, the borrowing constraint in cycle year $\tau - 2$ does not bind, but the household nevertheless saves nothing. If μ is less than $\hat{\mu}^{\tau-1}$, then the borrowing constraint in cycle year $\tau - 2$ binds; $A^1 = A^2 = \dots = A^{\tau-1} = 0$, and $A^\tau > 0$. If instead $\mu > \hat{\mu}^{\tau-1}$, then $A^{\tau-1} > 0$. Applying this reasoning to higher and higher values of μ yields the following result.

Proposition 1 *There exist positive and finite threshold values of μ , $\hat{\mu}^2 > \hat{\mu}^3 > \dots > \hat{\mu}^\tau$, such that $A^\kappa > 0$ if and only if $\mu > \hat{\mu}^\kappa$.*

Note that progressively higher values of μ generate positive assets for year $\tau - 1$ first, then for year $\tau - 2$, and so on backwards until year 1 of the cycle. The constraint always binds in the cycle's final year, so that $A^1 = 0$.¹⁰ We conclude that the borrowing constraint “switches off” at most once during the cycle. It switches back on in the year of the special expenditure. We use this result to link the level of assets to the stage in the cycle.

Proposition 2 *Assume that the constraint switches off in year κ of the deterministic cycle, where $1 \leq \kappa < \tau$. Then, because $W - T \geq C^\kappa > C^{\kappa+1} > \dots > C^{\tau-1}$, we have that $A^{\kappa+1} < A^{\kappa+2} < \dots < A^\tau$.*

In words, the saving towards the next periodic expenditure monotonically increases the level of assets.

¹⁰The borrowing constraint must bind at some point along any deterministic path, not just one which forms a nonstochastic cycle. Assume otherwise, so that for some t $A_{t'} > 0$ for all $t' \geq t$. This cannot be the optimal behavior, because increasing C_t , and hence all subsequent consumption levels, at the expense of reducing all future asset levels, gives higher present value utility.

3.3 Shortening of the Planning Horizon and the MPC

Zeldes (1984) noted that a binding borrowing constraint in the future works like a terminal condition which shortens the effective planning horizon. The household's response to an unanticipated temporary increase in $W_t - T_t$ on the nonstochastic cycle illustrates this. If the borrowing constraint binds in the year of the increase, then $MPC = 1$ as expected. If instead, the borrowing constraint is slack then, the household allocates the increase in current income across consumption between the present year in the cycle, κ , and the next time the borrowing constraint binds. The resulting marginal propensity to consume is

$$MPC^\kappa = \left(\frac{1 - \beta^{\tau-\kappa+1}}{1 - \beta} + \mu\beta^{\tau-\kappa} \right)^{-1}.$$

This exceeds the marginal propensity to consume of an unconstrained household facing the interest rate $\beta^{-1} (1 - \beta)$ if and only if $\mu < \beta/(1 - \beta)$.¹¹ The model's calibration satisfies this condition comfortably, so we proceed under this assumption.

We began this paper highlighting the empirical puzzle of MPC s substantially larger than the rate of interest for households with wealth. To see our model's implications for these observations, we differentiate MPC^κ above with respect to κ . The upper bound for μ signs the derivative positively. Therefore, we conclude:

Proposition 3 *If $\mu < \beta/(1 - \beta)$, and if the borrowing constraint becomes slack in year κ of the cycle, then $MPC^\kappa < MPC^{\kappa+1} < \dots < MPC^{\tau-1}$.*

Propositions 2 and 3 together imply that if we sampled households uniformly distributed across the deterministic cycle, we would find that MPC_t covaries positively with A_t among households with assets.

3.4 Persistence of Consumption Growth

We now turn to the model's implications for the persistence and volatility of consumption growth. Raising income permanently by one percent increases the household's assets permanently by the same percentage. This long-run connection between assets and income is absent from the standard model of an unconstrained household.¹² This connection generates

¹¹If instead, $\mu > \beta/(1 - \beta)$, then the periodic expenditure has such a large share that most of the temporary increase goes towards it.

¹²Consumption of a standard permanent-income household adjusts immediately and fully to an unexpected permanent wage increase, leaving assets unchanged.

persistence of aggregate consumption growth, because the household accumulates the new target asset level gradually.

To see this, we consider the consumption of τ households—each of which begins from a different point on the nonstochastic cycle—following an unexpected and permanent one percent increase in $W - T$. Since we are interested in the consumption decisions of households with positive assets, we proceed with the assumption that $\mu > \hat{\mu}^\tau$. We begin with the response of a household which just undertook its periodic expenditure, and hence has currently no assets. An unexpected and permanent wage increase of one percent raises the value of all resources available before the next periodic consumption by one percent. The homotheticity of the household’s preferences requires that the consumption of all goods also rises by one percent. This implies that its asset holdings at every date also rise, due to higher saving for the next periodic expenditure. For this household, the adjustment of assets and consumption is immediate because the entire effective planning horizon is still ahead. Therefore, consumption growth of a household beginning with no assets displays no persistence.

In contrast, if the household has positive assets, the shock raises the value of total resources available before the next periodic expenditure by less than one percent if the household has assets, so consumption between the period of the shock and the next periodic expenditure goes up by less than one percent. The household completes its consumption adjustment only after the next periodic expenditure.

Aggregating consumption responses of households with assets with those from households without assets yields an aggregate consumption series displaying persistence following an unexpected wage change. If wage innovations hit these households repeatedly, we would also find that their aggregate consumption growth has a smaller variance than does income growth. This follows automatically from the representation of consumption and income growth as moving averages of current and (for consumption) lagged income shocks. The variance of either series equals the sum of the squared moving average coefficients, and Jensen’s inequality implies that this is maximized by concentrating all of the response in one period, which is exactly what the moving average for income does by assumption. Since the response of consumption growth is more gradual, its unconditional variance is lower.¹³

¹³In [Galí’s \(1990\)](#) infinite-horizon framework, retirement savings mitigate the wealth effect of an income innovation and thereby reduce consumption volatility. This qualitatively resembles the effects of term saving on consumption volatility here, but [Galí](#) finds that the variance of aggregate consumption growth still exceeds that of income growth in his calibrated model.

3.5 The Model with Durable Goods

A fixed constraint on net assets like $A_{t+1} \geq 0$ has tractability but lacks realism. Middle-class households often carry substantial debts, and the vast majority of these are collateralized by durable goods.¹⁴ To better equip the model for quantitative work, we extend it with durable goods—which we think of as including housing—and collateralized debt.

Denote the stock of durable goods at the beginning of the year with S_t . We assume that the corresponding service flow is proportional to this stock, which depreciates at the constant rate δ . A simple specification for collateral value is

$$V_t = (1 - \pi) S_t, \quad (11)$$

where π is a required equity share—an exogenous parameter—corresponding to the down-payment rate. Limiting borrowing to the value of collateral yields

$$A_{t+1} \geq -V_{t+1}. \quad (12)$$

The definition of V_t in (11) implies that the good’s collateral value depreciates at the same rate as its service flow, δ . We can use this to express the accumulation of collateral value with the standard perpetual-inventory form:

$$V_{t+1} = (1 - \delta) V_t + (1 - \pi) (S_{t+1} - (1 - \delta) S_t). \quad (13)$$

If the borrowing constraint always binds, then (13) requires that collateralized debt be amortized at the physical depreciation rate. In reality, amortized mortgages and typical automobile loans require the borrower to repay the loan faster than this, so that the borrower’s equity share rises over time. We embody this into the model by supposing that a good’s collateral value depreciates faster than the flow of services it generates. This replaces the first instance of δ in (13) with the depreciation rate of a durable’s collateral value, ϕ .

$$V_{t+1} = (1 - \phi) V_t + (1 - \pi) (S_{t+1} - (1 - \delta) S_t). \quad (14)$$

This perpetual-inventory collateral accumulation equation reduces to (11) if $\phi = \delta$. When $\phi > \delta$, the required equity share for a given purchase increases as the good ages.

We now complete the extension by adapting the household’s preferences and budget constraint to the presence of durable goods.

¹⁴ [Aizcorbe, Kennickell, and Moore \(2003\)](#) use the 2001 Survey of Consumer Finances and estimate that 90 percent of all household debt is collateralized by homes and vehicles.

$$\sum_{t=0}^{\infty} \beta^t \{ (1 - \theta) \ln C_t + \theta \ln S_t + \mu_t \ln M_t \}, \quad 0 < \beta < 1, \quad 0 < \theta < 1, \quad (15)$$

$$C_t = W_t - T_t + RA_t + (1 - \delta) S_t - A_{t+1} - S_{t+1} - M_t. \quad (16)$$

The household's problem is to choose sequences of C_t , M_t , V_{t+1} , A_{t+1} , and S_{t+1} to maximize utility subject to the sequences of budget, borrowing, and collateral accumulation constraints, and the initial stocks, V_0 , A_0 and S_0 . We continue to denote the Lagrange multipliers on the budget and borrowing constraints with Ψ_t and Γ_t , and denote the multiplier on the collateral accumulation constraint with Ξ_t . After the obvious change to the expression for the marginal utility of ordinary non-durable consumption, the first-order conditions from the basic model apply to this extension. The additional conditions for the optimal choices of S_{t+1} and V_{t+1} are

$$\Xi_t = \Gamma_t + \beta(1 - \phi)\Xi_{t+1}, \quad (17)$$

$$\Psi_t - \Xi_t(1 - \pi) = \beta \frac{\theta}{S_{t+1}} + \beta(1 - \delta)(\Psi_{t+1} - \Xi_{t+1}(1 - \pi)). \quad (18)$$

Both of these equations have straightforward interpretations. Iterating (17) forward expresses Ξ_t as the utility value of relaxing the present and future borrowing constraints by marginally increasing collateral value. Since Γ_{t+j} must exceed zero for some $j \geq 0$, Ξ_t is positive. Its magnitude summarizes the present importance of anticipated borrowing constraints. Hence, Ξ_t expresses the degree to which the household is globally constrained.

If we artificially set Ξ_t and Ξ_{t+1} to zero, then (18) reduces to the standard first-order condition for optimal investment in durable goods. This reveals that this household subtracts the marginal utility from the accompanying expansion of collateral, $\Xi_t(1 - \pi)$ from the utility cost of durable goods consumption in year t .

3.6 A Life-Cycle Interpretation of the Model

We have motivated our infinite-horizon model's periodic expenditure as reflecting life-cycle events. We give this motivation an explicit foundation in the Appendix which presents a life-cycle overlapping generations model of a dynasty with altruistic members. That model

features tax-advantaged Individual Retirement Accounts. We show that it is identical to the infinite horizon model around its deterministic cycle if the dynasty fully exploits its IRA investment opportunities in spite of its impatience. Here we discuss the model's setup and summarize the results of its analysis. This discussion omits durable goods and includes only the standard borrowing constraint that allows no debt. The Appendix discusses the natural extension with durable goods and collateralized household debt.

The model of the Appendix consists of a dynasty with γ overlapping generations. Every τ years a new member is born, and he dies $\tau\gamma$ years later. In the first $\tau(\gamma - 1)$ years of life, an individual earns an annual after-tax labor income of $W_t N$ by inelastically supplying N units of labor per year. In the last τ years, the individual is retired and earns nothing. All members pay annual lump-sum taxes T_t . Member i derives utility from ordinary nondurable consumption, C_t^i , and end-of-life care, M_t^i . From these he gets the utility flow

$$\ln C_t^i + \mu_t^i \ln M_t^i,$$

where $\mu_t^i = \mu > 0$ in the final year of life and $\mu_t^i = 0$ in other years. We call M_t^i “end-of-life care” only for the sake of concreteness. The analysis of this model would be identical if we instead supposed that M_t^i represents some other life-cycle expenditure, such as education in the first period of life. Dynasty members discount future utility at the common rate $0 < \beta < 1$, and they are perfectly altruistic towards the dynasty's other born and unborn individuals.

The dynasty can save either by accumulating ordinary bonds or by investing in its working members' Individual Retirement Accounts. The bond's rate of return equals $R < 1/\beta$, so the dynasty is impatient in the usual sense. The United States exempts the capital income of IRA investments from taxation, so we assume that an IRA's rate of return in the dynastic model satisfies $R^* > R$. As in the United States, each individual's annual contributions to the IRA may not exceed a maximum, and each retiree's annual withdrawals must exceed a minimum fraction of the remaining balance.

Dynastic optimization requires equalization of all members' marginal utilities from ordinary consumption ($C_t^i = C_t^{i'}$) and equalization of marginal utility from ordinary consumption and end-of-life care for the oldest living member ($M_t^i = C_t^i \mu_t^i$). In the Appendix, we use these conditions to express the dynasty's utility function (which sums the utilities of all born and unborn members) as the utility function in (1). We then show that if β is not too low and the shocks to earnings are small, then the IRA's tax advantage generates a corner solution: All working individuals make the maximum contribution, and all retirees make the minimum

withdrawal. Specifically, obtaining this corner solution requires $R^*\beta > 1$.¹⁵ With this result in hand, we can define the dynasty’s net labor income as the earnings of its working members net of their IRA contributions plus the retired member’s IRA withdrawal. Since the dynasty’s utility function and budget constraints are identical to those of the infinitely lived household, the dynastic model’s solution corresponds exactly to the model in this section.

We also assume that the IRA contribution ceiling is such that the only transfers across dynasty members are from workers to retirees to finance their share in the end-of-life care. In this case, the ratio of nonretirement assets to disposable income is equalized across all members of the dynasty. This is realistic in the presence of a pay-as-you-go public pension, which, for simplicity we omit from the model. For working members, disposable income is labor income net of taxes and pension funds contributions, and for retirees it is pension income. This equalization will be relevant for the calibration of the model, because the ratio of nonretirement assets of working age individuals to their disposable labor income is interpreted as reflecting the dynasty’s optimal decisions.

4 Quantitative Analysis

In this section, we calculate the calibrated model’s responses to transitory income changes and balanced-budget tax experiments and then discuss the evidence in Section 2.1 from the model’s perspective. Comparing the results with the evidence requires us to resolve a financial indeterminacy: Since the household pays the same interest rate on debt as it receives on savings, it can save either by purchasing assets or by accumulating equity in its durable goods. We resolve this by assuming that repaying debts faster than required and then extracting the funds with new borrowing incurs a small cost. To avoid it, the household repays its debt at the minimum required pace. Under this assumption, we measure the household’s gross debt with V_t and its gross assets equal $A_t + V_t \equiv A_t^g$.

As we noted above, the periodic expenditure motivates an impatient household to save. Accordingly, we choose the parameters governing saving (μ and τ) using observations of middle class households’ financial assets from the 2001 Survey of Consumer Finances. After introducing that evidence, this section presents the model’s calibrated parameters and discusses its nonstochastic cycle. It then proceeds with the calculation of households’ responses

¹⁵We require $R^*\beta > 1$ to obtain a corner solution, but this is not sufficient when $\tau > 1$. In the Appendix we derive a condition on β , R , and R^* sufficient for a corner solution in the dynastic model’s deterministic cycle, equation (A14).

to a temporary wage increase and to several intertemporally-balanced tax reductions. We compare the latter to the evidence from the 2001 tax cuts surveyed above. We also use the model to quantitatively interpret excess smoothness and persistence in aggregate consumption.

4.1 Middle Class Financial Assets

The evidence on the marginal propensity to consume surveyed above comes primarily from the tax rebates of 2001, so we draw our observations from the 2001 Survey of Consumer Finances. For each surveyed household, the SCF contains the values of several financial assets (including IRA accounts not associated with an employer) as of the interview date. It also reports the year 2000 pre-tax labor income of both the respondent and her or his spouse, contributions to employer-sponsored retirement savings programs, pre-tax capital income, and the household's 2000 tax year Adjusted Gross Income.

In the model, the consumption and saving decisions are homogeneous of degree one in net labor income. Hence, we consider the financial assets of each household relative to its disposable labor income. For the construction of both variables in this ratio, we follow the dynastic interpretation of our model in Section 3.6 and the Appendix. The empirical counterpart of A_t^g is financial assets excluding balances in tax-advantaged retirement accounts: stocks, bonds, as well as balances in checking, savings, money market, and mutual fund accounts. For the measurement of disposable labor income, we first calculate the household's income taxes paid by applying the reported AGI to the year 2000 tax table. We then apportion this between labor and capital using reported capital income. To get our measure of after-tax labor income, we subtract the resulting labor income taxes, the household's FICA and Medicare taxes, contributions to employer-sponsored retirement plans, and IRA contributions from the household's wages, salaries, and self-employment earnings.¹⁶

We intend our model household to represent the working middle class, so we apply several screens to the data. We first keep only households with heads between 25 and 64 years old with positive after-tax labor income. We remove the rich by dropping the wealthiest five percent of the remaining households, and we remove the poor by deleting any household that received Unemployment Insurance Benefits, Food Stamps, or TANF payments in the previous year. The initial sample represents 106 million households. Table 2 reports the

¹⁶The 2001 SCF has no information on year 2000 IRA contributions, but it does list all of the household's IRA accounts. Our calculations suppose that each individual with an IRA made the maximum legal contribution (\$2,000).

Households Represented in Original Sample,	106,493,608
& with heads between 25 and 64 years old,	78,079,336
& with positive after-tax labor income,	68,659,384
& are among least wealthy 95% of remaining households,	65,232,296
& that received no UI, Food Stamps, or TANF,	59,728,852

Table 2: Number of Households Represented in the 2001 SCF

number of households remaining after each screen. The final sample represents 60 million households.

Table 3 reports the income-weighted average ratios of wealth to disposable labor income. The leftmost cell’s estimate uses the entire screened sample. The other cells report the average ratios by deciles of this ratio. For the complete sample, the average ratio is 54.9 percent. This greatly exceeds the ratio for the fifth decile, 12.1 percent, which approximates the distribution’s median. Apparently wealth is concentrated in this sample in spite of omitting the very wealthy and poor. Qualitatively, the model’s deterministic cycle can generate such concentration because the household saves *at an increasing rate* for the next periodic expenditure. We will use the overall average ratio of assets to disposable labor income to calibrate the model, and then compare in Section 4.3 the resulting model’s distribution of wealth with the evidence presented here.¹⁷

4.2 Calibration

We now proceed to calibrate the model’s parameters. For this we combine the calibration strategy from Campbell and Hercowitz (2009) with the wealth observations from Table 3 and considerations from the life-cycle interpretation of the model. We interpret each of the model’s units of time as one year, and set $R = 1.04$. The physical depreciation rate δ is set at the stock-weighted average of those for vehicles and residences from the Bureau of Economic Analysis Fixed Reproducible Tangible Wealth, 0.04.

¹⁷In the simplest overlapping generations model without altruism, households accumulate wealth as they age to fund consumption in retirement. Thus, one might suspect that the wealth heterogeneity in Table 3 reflects age heterogeneity. We have investigated this possibility by regressing the wealth to labor-income ratio on the household head’s age and age squared. Although the estimated regression indicates that the ratio does rise with age, its R^2 is less than 1 percent. Thus, the data reveal substantial wealth heterogeneity across middle-class households of the same age. It is this heterogeneity that our model addresses.

Full	Deciles									
Sample	1	2	3	4	5	6	7	8	9	10
54.9	0.2	1.9	4.2	7.3	12.1	19.4	31.3	55.0	111.7	336.5

Table 3: Ratios of Financial Assets to Annual Disposable Labor Income ($\times 100$)

Note: Each cell reports a weighted average of nonretirement financial assets to labor income net of federal income taxes, Social Security taxes, and contributions to tax-advantaged retirement accounts. The leftmost cell uses the entire sample, while the remaining cells use observations grouped by deciles of this ratio.

The extensive deregulation of the mortgage market in late 1970s and early 1980s changed the effective equity requirements for household debt. Since we wish to compare our model’s dynamics with observations from the 2001 tax cut, we calibrate π and ϕ using observations from loan contracts relevant for that year. The down-payment rate is a weighted average of those for vehicles and residences. For cars, the average loan-to-value ratio from Federal Reserve Statistical Release G19 over the 1995-2004 period is 0.92. For homes, existing data on first loan-to-value ratios over this period is not useful because “down payment assistance” loans can lower the effective equity share held at purchase. We use observations from the 2001 Survey of Consumer Finances from households who purchased homes in the 12 months preceding the interview. Their average equity share is 0.175. We expect households to accumulate very little additional equity in the first year of home ownership, so we take this as a measure of the effective down-payment rate. The relevant weighted average of these two down payment rates is $\pi = 0.108$.¹⁸ We set ϕ to 0.074, four times the quarterly rate of 0.0186 used in [Campbell and Hercowitz \(2009\)](#).

The four remaining parameters are τ , μ , β , and θ . For guidance calibrating the first three, we draw upon the life-cycle version of the model as discussed in Section 3.6 and the Appendix. As mentioned there, the IRA tax advantage generates corner solutions for IRA contributions and withdrawals if impatience is not too strong. Specifically, the requirement is that $R^*\beta > 1$. Applying a marginal tax rate of 30 percent to $R = 1.04$ yields $R^* = 1 + (R - 1)/(1 - 0.3) = 1.0595$. We choose $\beta = 1/1.055$, which satisfies this condition and exceeds the interest rate R by one and a half percentage points.¹⁹

¹⁸See Appendix A of [Campbell and Hercowitz \(2009\)](#) for details on the calculation of π .

¹⁹Given $R = 1.04$ and $R^* = 1.0595$, our choice of $\beta = 1/1.055$ satisfies $R^*\beta > 1$ and the more stringent sufficient condition in equation (A14). We note here that the actual value of R^* has no bearing on our analysis beyond the discipline it provides in choosing β .

β	R^{-1}	θ	τ	μ	π	ϕ	δ
0.95	0.96	0.42	10	1.82	0.108	0.074	0.040

Table 4: Calibrated Parameter Values

In the dynastic model, with the suitably chosen IRA contribution ceiling, ordinary saving finances only the periodic expenditure. Therefore, observations of nonretirement assets of working-age individuals can be used to infer the size of the periodic expenditures (determined by μ) and the number of years between them (τ). In practice, we arbitrarily set τ to ten years, although the results are nearly the same with nine and eleven years. We then choose μ to match the income-weighted average ratio of nonretirement financial wealth to disposable labor income from Table 3, 54.9 percent. We select θ to match the share of total personal consumption expenditures accounted for by nondurable goods and nonhousing services. Over the 1983–2006 period, this equaled 21.1 percent. Our choice of μ depends on that for θ , and vice versa. Therefore, calibration requires the straightforward solution of two nonlinear equations in two unknowns. Table 4 lists the resulting parameter values.

4.3 The Deterministic Cycle

With the calibrated parameter values in hand, we calculate the extended model’s deterministic cycle numerically. Figure 1 plots the household’s consumption choices, assets, and Lagrange multipliers along the calibrated model’s computed nonstochastic cycle. We place the periodic expenditure in the cycle’s final year and normalize $W - T$ to one.

The liquidity constraint binds in only one of the cycle’s years, that of the periodic expenditure. Otherwise, the Euler equation holds and $\Gamma_t = 0$. Since Ξ_t equals a present discounted value of current and future values of Γ_t , it always exceeds zero. It grows as the periodic expenditure and the accompanying binding borrowing constraint approach. For this reason, we measure the degree to which the household is globally liquidity constrained at any date with Ξ_t . During the years in which the Euler equation holds it requires both C_t and S_{t+1} to decline at the rate $1/\beta R - 1$.

The lower right panel shows A^g . As in the basic model without durable goods, gross financial assets grow at an increasing rate as the household approaches the periodic expenditure. If individuals are uniformly distributed over the expenditure cycle, a sample drawn from this model at the beginning of the year would have average ratios of assets to net labor income by decile equal to the values in this panel. Hence, it is the model’s counterpart to

Figure 1: The Calibrated Model's Nonstochastic Cycle

Table 3. In the data, the ratio of the fifth decile’s average to the mean is 0.22, while in the model, this ratio is 0.69. Hence, inequality in the model is significantly lower than in the middle-class sample for that table.

4.4 Dynamic Responses to a Temporary Income Change

Figure 2 shows the responses of nondurable consumption to a transitory, unexpected, and marginal transfer of labor income. Each panel plots a different timing of the transfer, from $\kappa = 1$, i.e., one year after the expenditure, to $\kappa = 10$, the year of the expenditure. The response is expressed as a percentage of the transfer, so its value in the year of the transfer is a partial marginal propensity to consume (which only reflects nondurable purchases).

The highest response is 78 percent corresponding to $\kappa = 10$. In this year the borrowing constraint binds, which explains the high response. When the transfer is received in years 1 to 9, i.e., when the borrowing constraint does not bind, the responses are smaller and range from 6 percent to 14 percent. The interesting feature of these responses is that they increase as the next periodic expenditure gets closer. This happens simultaneously with increasing asset levels. Hence, the conclusion of Proposition 3 holds good for this calibration of the extended model.

Figure 2 also shows that the household smooths consumption over a short horizon that extends only until the next periodic expenditure. For example, when $\kappa = 5$, the response is smooth at the 7 percent level for four years. Seven years after the periodic expenditure, consumption falls back to its initial level. We demonstrated analytically the shortened planning horizon in the basic model without durable goods or collateralized debt. We see here that the binding borrowing constraint effectively reduces the planning horizon in spite of the intertemporal connections these changes introduced.

Figure 3 portrays the corresponding impulse responses for durable goods purchases. These responses are much larger than those for nondurable consumption because households use part of the windfall as down payments on durable goods. Credit finances the purchase’s balance. Aside from their far greater magnitudes, these initial responses resemble those of nondurable consumption. They increase from 55 percent when the expenditure is nine years away to 151 percent in the year before the expenditure. In the expenditure year (when the borrowing constraint binds) it equals 206 percent. The decreases of durable purchases in future years show that these increases are mostly changes in the timing of purchases that would have occurred before the next periodic expenditure without the windfall.²⁰

²⁰We remind the reader that these percentages are expressed relative to the temporary wage increase and

Figure 2: Responses of Nondurable Consumption to a Temporary Wage Increase

Each panel plots the responses of nondurable consumption over a 10-year horizon to a marginal increase in income lasting one year. All responses are scaled relative to the income change, so their values in year one can be interpreted as partial marginal propensities to consume. In the headings, κ refers to the year of the household's deterministic cycle in which the income change occurs. Tick marks indicate the household's initial response, its maximum response, and the maximum response across all households. The horizontal helper lines mark 20, 40, and 60 percent of the income increase. The income change occurs in year 1. Please see the text for further details.

Figure 3: Responses of Durable Purchases to a Temporary Wage Increase

Each panel plots the responses of durable consumption purchases over a 10-year horizon to a marginal increase in income lasting one year. All responses are scaled relative to the income change, so their values in year one can be interpreted as partial marginal propensities to consume. In the headings, κ refers to the year of the household's deterministic cycle in which the income change occurs. The periodic expenditure coincides with the temporary increase for $\kappa = 20$. Tick marks indicate the household's maximum and minimum responses as well as the maximum and minimum responses across all households. The income change occurs in year 1. Please see the text for further details.

4.5 Dynamic Responses to a Tax Cut

We now proceed to compute the responses to a tax cut, and compare the results to the effects of the 2001 tax cut reviewed in Section 2.1. The simulation is based on the following assumptions about the households' perceptions of fiscal policy: (a) The tax cut is unexpected when implemented, (b) the tax cut does not last very long, (c) the future path of government spending remains unchanged, and (d) the government increases future taxes permanently to pay the interest on the additional government debt incurred.

We consider three possibilities for the tax cut's duration, one, three, and five years. In all cases we keep the tax cut's per year magnitude the same, so the permanent tax increases that follow them increase with the tax cut's duration. Figures 4 and 5 are the counterparts to Figures 2 and 3 for the experiment with a one-year tax cut. As expected, raising future taxes to balance the government budget generally lowers the responses of both durable and nondurable goods. However, these changes are in general small. Hence, the consumption response following a temporary income windfall changes little if future after-tax income drops to balance the government budget. This reflects the shortened planning horizon generated by the anticipation of a binding borrowing constraint. In light of this result, Shapiro and Slemrod's puzzling finding that households' consumption responses to the 2001 tax rebates have no association with their stated expectations about future changes to government spending and taxes makes sense.

The experiment with a three-year-long tax cut (not shown) generates higher responses than the one-year tax cut in cycle years $1 \leq \kappa \leq 9$. This follows from consumption smoothing of a more prolonged tax cut over a short horizon. When the borrowing constraint does bind ($\kappa = 10$) the response is almost identical to the analogous response from the one-year tax cut. This is what the basic model leads us to expect. Then, when the household is currently constrained, future income has no influence on current consumption. The addition of durable goods and collateralized debt hardly changes this conclusion. The experiment with a five-year-long tax cut yields similar but quantitatively larger results for the unconstrained periods.

We now compare these results with the evidence on the tendency to spend most of the rebate by wealth groups presented by Shapiro and Slemrod.²¹ For this, we adopt their assumption of a positive link between the probability of a survey respondent declaring that

not relative to the initial durable goods stock. Therefore, the negative responses after the initial windfall reflect a decrease in durable purchases rather than an outright sale of durable goods.

²¹Recall from Section 2.1 that Johnson, Parker, and Souleles (2006) whose point estimate of the relationship between the marginal propensity to consume and household wealth was *U*-shaped. We interpret their findings as reinforcing those of Shapiro and Slemrod (2003). We focus on their results only for parsimony.

Figure 4: Responses of Nondurable Consumption to a One-Year Tax Cut

Each panel plots the responses of nondurable consumption over a 10-year horizon to a small tax cut lasting one year followed by a permanent tax increase that balances the government's intertemporal budget. All responses are scaled relative to the tax cut, so their values in year one can be interpreted as partial marginal propensities to consume. In the headings, κ refers to the year of the household's deterministic cycle in which the tax cut occurs. The periodic expenditure coincides with the temporary tax cut for $\kappa = 20$. Tick marks indicate the household's initial response, its maximum response, and the maximum response across all households. The horizontal helper lines mark 20, 40, and 60 percent of the income increase. The tax cut occurs in year 1. Please see the text for further details.

Figure 5: Responses of Durable Purchases to a One-Year Tax Cut

Each panel plots the responses of durable consumption purchases over a 10-year horizon to a small tax cut lasting one year followed by a permanent tax increase that balances the government's intertemporal budget. All responses are scaled relative to the tax cut, so their values in year one can be interpreted as partial marginal propensities to consume. In the headings, κ refers to the year of the household's deterministic cycle in which the tax cut occurs. The periodic expenditure coincides with the temporary tax cut for $\kappa = 20$. Tick marks indicate the household's minimum and maximum responses as well as the minimum and maximum responses across all households. The tax cut occurs in year 1. Please see the text for further details.

she would spend most of her rebate and her actual (unobserved) *MPC*. Assuming also that their typical survey respondent considers only the down payment for a durable good purchase as expenditure—and not the additional debt incurred—we multiply the response of durable good purchases by the required down payment rate ($\pi = 0.108$) before adding it to the contemporaneous response of non-durable consumption. We define the sum of these two components as the overall marginal propensity to consume (*MPC*), which we tabulate in Table 5 by the cycle year of the tax cut or transfer, along with that year’s asset level. We measure each household’s gross assets as the average of their values at the beginning and end of the year—as in reality assets may be adjusted continuously. The first column of Table 5 lists κ , the year of the household’s nonstochastic cycle in which the experiment begins. The next column presents gross assets, and the remaining columns give the *MPC*s from the experiments.

The most noticeable feature of Table 5 is the positive association between gross financial assets and the marginal propensity to consume. The relationship, however, is not monotonic; the highest *MPC* does not match the highest level of assets; it corresponds to $\tau = 10$, i.e., the year when the borrowing constraint binds.²² Hence, the *MPC*s decrease when assets exceed 0.7. Nevertheless, they exceed the *MPC*s for households with low assets. This resembles the link found by Shapiro and Slemrod, as reported in Table 1, for households with positive stocks. They explain this positive correlation as adjustment of assets to a target level: Individuals with low assets spend relatively little in order to build up their assets towards the target, while those with high assets do the opposite.

Shapiro and Slemrod’s explanation and ours contrast each other in an interesting way. In their target-assets argument, high assets imply *temporarily high* consumption because the household dissaves. Here, high assets imply *temporarily low* consumption; the borrowing constraint will bind soon, and consumption smoothing over the short remaining horizon generates low consumption in the present as well. If we model the target-asset behavior with a quadratic term on the deviation of assets from target, saving involves a positive marginal cost when assets exceed the target. Hence, when a consumer in that framework receives temporary income while having high assets, her consumption reacts strongly in spite of her low marginal utility, because saving is costly. The household in the present model has a strong consumption response because of high marginal utility. Since he is globally constrained, his response qualitatively resembles that of a currently borrowing constrained individual.

²²For $\tau = 10$, assets are not the highest in Table 5 because they average the highest beginning-of-period level A^τ with the lowest end-of-period level $A^1 = 0$.

κ	Gross	One Year	Balanced-Budget Tax Cut for		
	Assets ⁽ⁱ⁾	Transfer	One Year	Three Years	Five Years
1	0.03	12	8	25	44
2	0.10	12	9	28	48
3	0.20	13	10	31	54
4	0.32	14	11	35	61
5	0.46	16	13	40	70
6	0.63	18	15	47	82
7	0.82	20	18	57	77
8	1.04	24	22	70	70
9	1.29	30	29	59	59
10	0.72	100	100	100	100

Table 5: The *MPC* and Financial Assets

Note: (i) The measure of Gross Assets in year t in this column equals $(RA_t^g + A_{t+1}^g)/2$.

As shown in Table 1, [Shapiro and Slemrod](#) find a large fraction of respondents with zero stocks-market assets (43 percent), whose tendency to spend is higher than those with low stocks levels. They interpret this group as non-savers, i.e., individuals who are permanently borrowing constrained. Households with zero assets are absent from Table 5. This is consistent with our interpretation of the model as reflecting middle class households, which hold at least a small amounts of assets.²³

4.6 Interpreting Evidence on Consumption Growth Stickiness

We now turn to the calibrated model's quantitative implications for excess smoothness and persistence of consumption growth. Since the time period in our model is one year, we begin with a test of excess smoothness similar to [Campbell and Deaton's \(1989\)](#) quarterly analysis

²³This is a result of the chosen calibration and not an intrinsic feature of the model. Increasing τ to 15, i.e., extending the cycle and correspondingly recalibrating μ , makes the borrowing constraint bind not only in the expenditure year, but also in the following one, $\kappa = 1$. In this case, therefore, households have zero assets at the beginning and end of cycle year $\kappa = 1$, while the corresponding *MPC* is high because then the borrowing constraint binds. If we increase τ further to 20, then the borrowing constraint binds for the first six years after the periodic expenditure. Hence, this model can generate a fraction of households with zero assets and high *MPC*. However, these households would have zero assets only temporarily. They would not constitute a separate group of permanently low-wealth households.

using annual data. Their procedure is based on the permanent-income consumption equation

$$C_t = \frac{\vartheta r}{1+r} \left[A_t + \sum_{i=0}^{\infty} (1+r)^{-i} E_t Y_{t+i} \right],$$

where r is the real interest rate, ϑ is the constant ratio of nondurable consumption to total consumption (assuming a constant relative price), A_t is current assets and Y_t is disposable labor income. This follows a random walk with drift: $\ln(Y_t/Y_{t-1}) = \mu + \varepsilon_t$, $\varepsilon_t \sim N(0, \sigma^2)$. All variables are measured at the household level, so μ equals income growth per capita. This income process implies²⁴

$$(1+r)^{-i} E_t [Y_{t+i}] = Y_t (1+r)^{-i} \exp(i(\mu + \sigma^2/2)).$$

Using this, the consumption equation can be expressed as

$$C_t = \frac{\vartheta r}{1+r} \left[A_t + Y_t \sum_{i=0}^{\infty} \frac{\exp(i(\mu + \sigma^2/2))}{(1+r)^i} \right],$$

or

$$C_t = \frac{\vartheta r}{1+r} A_t + Y_t \left(\frac{\vartheta r}{1+r - \exp(\mu + \sigma^2/2)} \right).$$

From here it follows that the percentage change in consumption after a permanent income innovation equals

$$\frac{\Delta C_t}{C_t} = \frac{\vartheta r}{1+r} \frac{\mu A_0}{C_t} + \left(\frac{\vartheta r}{1+r - \exp(\mu + \sigma^2/2)} \right) \frac{\Delta Y_t}{C_t}, \quad (19)$$

where Δ denotes the first-difference operator, and $\Delta A_t = \mu A_0$ for all t . The latter holds because income innovations change consumption by exactly the same amount, without altering the growth of assets. Hence, the standard deviation of consumption growth according to the permanent-income hypothesis is

$$SD^{\text{pih}} \left(\frac{\Delta C_t}{C_t} \right) \approx \left(\frac{\vartheta r}{1+r - \exp(\mu + \sigma^2/2)} \right) SD \left(\frac{\Delta Y_t}{C_t} \right). \quad (20)$$

The approximation in (20) ignores the term $\mu A_0/C_t$, whose coefficient, when setting a period to one year, is small compared with the coefficient of the income change term. Except for

²⁴This expectation is finite if $\exp(\mu + \sigma^2/2) < 1+r$.

setting $r = 0.04$, we estimate the other terms in this equation using the sample 1983 – 2007.²⁵ The data are per capita aggregate disposable labor income and nondurable consumption less housing services. Excess smoothness is observed if $SD(\Delta C_t/C_t) < SD^{\text{pih}}(\Delta C_t/C_t)$.

The necessary estimates for this comparison are $\mu = 0.017$, $\sigma = 0.016$, $\vartheta = 0.723$, $SD(\Delta Y_t/C_t) = 0.016$, and $SD(\Delta C_t/C_t) = 0.008$. The result is $SD^{\text{pih}}(\Delta C_t/C_t) = 0.020$, which is much larger than the standard deviation of actual consumption growth. Hence, the “excess smoothness” stressed by Deaton (1987) and Campbell and Deaton (1989) appears strongly in this sample. Note also that the volatility of income, σ , is much higher than the volatility of consumption.

The sample displays also “consumption growth stickiness”. The autocorrelation coefficient of consumption growth is 0.40 with a t -statistic of 2.36. This coefficient is lower than the annualized estimate in Carroll, Slacalek, and Sommer (2009) of 0.62, although it is not far considering the estimate’s standard error, 0.17.

We noted in Section 3.4 that the simple version of the model can qualitatively reproduce the excess smoothness and persistence of the data. To examine the calibrated model’s quantitative performance on this dimension, we calculate the response of aggregate consumption following a permanent one percent wage increase for a population of ten households evenly distributed across points in the nonstochastic cycle. We use this impulse response as the infinite moving-average representation of aggregate consumption in the face of a unit-root income process with a one-percent standard deviation. Hence, for this calculation, $\sigma = SD^{\text{pih}}(\Delta C_t/C_t) = 0.01$. With the calibrated parameters, $SD(\Delta C_t/C_t) = 0.0067$. Thus, the model generates substantial excess smoothness. This occurs because aggregate consumption responds initially by only 0.67 percent, and it then completes its long-run one-percent adjustment slowly. This response generates also a small positive autocorrelation of consumption growth, 0.11. Overall, we find that liquidity constraints of households holding assets can contribute to the resolution of the consumption smoothness and persistence puzzles.

5 Concluding Remarks

How liquidity constrained are middle-class households in the U.S.? To address this question, we developed a model where households are home owners and hold financial assets, and mea-

²⁵This sample is chosen because of the drastic deregulation of the mortgage market in 1982 and the dramatic macroeconomic events that began in August 2007, which affected substantially the financial market and hence the environment for dynamic optimization.

sured liquidity constraints with the fraction spent out of a temporary tax rebate—compared to the unconstrained Ricardian response of zero.

In the model, a future binding borrowing constraint effectively shortens the planning horizon of households who are infinite-horizon planners. These households value liquidity in spite of being currently unconstrained. This model has two main implications:

- The spending responses to a transitory transfer in the model are much higher than for a permanent-income consumer. The responses to a transitory transfer and to a tax cut financed by a future permanent tax hike are very similar. In other words, future tax changes have little effect on current decisions. This implies that the response of these households to a balanced-budget tax rebate differs greatly from the zero-response of a Ricardian permanent-income consumer.
- The volume of assets owned reflects a forthcoming demand for liquidity rather than a liquidity surplus arising from past luck. This feature generates a positive relationship between assets and the marginal propensity to consume out of temporary income.

The second implication provides a rationalization of the finding in [Shapiro and Slemrod \(2003\)](#), that among households with positive amounts of shares, the fraction of households who spent most of the 2001 rebate increases with stock ownership.

In the calibrated version of the model, the average *MPC* from a one-year tax cut is 24 percent. For a five-year tax cut, it equals 66 percent. These figures are realistic given the evidence on the 2001 tax rebate. Our interpretation of these results is that middle-class households face quantitatively significant liquidity constraints.

Appendix: A Dynastic Foundation for the Model

This Appendix presents a dynastic overlapping generations model with perfect altruism and tax-advantaged individual retirement accounts. The model's IRAs carry a rate of return high enough to induce the dynasty's working members to save in them at the maximum allowed rate and withdraw from them as slowly as legally possible. Each individual gains utility from medical care at the end of life, which is a periodic expenditure for the dynasty. With these preferences and savings opportunities, there is a deterministic cycle which is equivalent to that of the infinitely lived household considered in the text, with income redefined as wages less IRA contributions plus IRA withdrawals.

The calibration of our model considered the ratio of the nonretirement financial assets of working-age individuals to their disposable labor income net not only of taxes but also of IRA contributions. We justify this strategy here using the dynastic model. For this, we represent the dynasty's decisions by giving each member a separate ordinary savings account. One implementation of the dynasty's consumption and savings plan has all dynasty members maintaining identical balances in them. With the appropriate IRA contribution limit, the only transfer payments required to equalize consumption are from working members to the retiree to fund end-of-life medical care. In this case, the dynasty's ratio of ordinary savings to current disposable income equals any working member's ratio of ordinary savings to wages net of IRA contributions. Furthermore, the corner solution for IRA contributions and withdrawals implies that the dynasty's marginal propensity to consume out of a given income and tax shock equals the corresponding value from the model in the text.

A.1 Tax-Advantaged Retirement Savings in the U.S.

Current U.S. tax law provides two kinds of tax-advantaged vehicles for retirement savings. The first is the *pre-tax contribution* account. The leading examples are the IRA account, its 401k and 403b variants, and numerous employer-sponsored plans. Generally, individuals (and possibly their employers) may make limited contributions to their accounts' balances. For example, in 2009, an individual may contribute up to \$5,000 per year (\$6,000 for those 50 and older) to an IRA. In most cases, these contributions are deductible from the individual's current taxable income. Realized capital gains, dividends, and interest accumulate in the account tax free thereafter. The IRS heavily penalizes withdrawals from these accounts before the individual turns $59\frac{1}{2}$. After this date, it taxes withdrawals as ordinary income. The IRS also imposes minimum withdrawals on those $70\frac{1}{2}$ or older. These approximately

empty the account over the individual's expected remaining life.

The second kind of tax-advantaged retirement savings account, the Roth IRA, differs from the traditional IRA in two respects. First, contributions to a Roth IRA account's balance are *not* deductible from current income. Second, the IRS does *not* count withdrawals from Roth IRA accounts towards current income. Contributions to a Roth IRA reduce dollar for dollar the maximum contribution to a traditional IRA.²⁶

The financial benefit from these retirement accounts comes from their favorable treatment of interest, dividends, and capital gains. To see this, it is helpful to note that funding either kind of account with the same consumption reduction yields the same after-tax resources in retirement if the marginal tax rate is constant. Suppose an individual reduces current consumption by \$1 to purchase a security with gross return R^* in the Roth IRA. In retirement, the assets are $\$R^*$. If the individual faces a tax rate of λ both while working and in retirement, investing $\$1/(1 - \lambda)$ in a traditional IRA with the same gross return requires the same \$1 reduction in current consumption and yields the same after-tax payout $\$(1 - \lambda) \times R/(1 - \lambda) = \R^* . An after-tax dollar invested outside of a tax-advantaged account yields only $\$(1 + (R^* - 1) \times (1 - \lambda)) < \R^* in retirement, because the IRS taxes that capital income. Since the two IRA accounts manifest their benefits in a higher after-tax return on investment, the model's IRA has this as its only advantage.

A.2 A Dynastic Model

Consider a dynasty of individuals with names $i = 0, 1, 2, \dots$. Individual i is born in year τi and lives for $\tau\gamma$ years. Hence, a new member is born every τ years and there are γ overlapping living generations. Note that older individuals have smaller names, as with kings: Henry IV preceded Henry V who in turn preceded Henry VI. An individual provides a labor supply N in the first $(\gamma - 1)\tau$ years of life and is retired in the last τ years. Each individual derives utility from ordinary nondurable consumption and end-of-life care. We denote individual i 's consumption of these goods in year t with C_t^i and M_t^i . Utility from this consumption is

$$\ln C_t^i + \mu_t^i \ln M_t^i,$$

where $\mu_t^i = \mu > 0$ when $t = (i + \gamma)\tau$ —the final year of life—and $\mu_t^i = 0$ at other times.

All dynasty members discount future utility at the common rate $0 < \beta < 1$, and they are perfectly altruistic towards the dynasty's other members both born and unborn. Therefore, they all rank consumption streams with the same utility function. We denote the name of

²⁶Information about tax-advantaged retirement saving in the U.S. is available in [IRS Publication 590](#).

the youngest living dynasty member with $\bar{i}(t)$. The name of the oldest living member is $\underline{i}(t) \equiv \bar{i}(t) - \gamma + 1$. The dynasty's welfare function is

$$\sum_{t=0}^{\infty} \beta^t \sum_{i=\underline{i}(t)}^{\bar{i}(t)} (\ln C_t^i + \mu_t^i \ln M_t^i). \quad (\text{A1})$$

The dynasty members earn the wage W_t for each unit of time sold in the labor market. They also have access to a bond market. They can purchase bonds through either an ordinary savings or an individual retirement account. Their annual gross interest rates paid are $R < 1/\beta$ and $R^* > R$. The first inequality implies impatience, as in the text, and the wedge between the two accounts' rates reflects the exemption of the IRA's capital income from taxation.

We denote the IRA balance of individual i at the end of year $t - 1$ with B_t^i . The contributions of i while working are I_t^i , and the withdrawals of i while retired are J_t^i . Therefore, the accumulation equations

$$B_{t+1}^i = R^* B_t^i + I_t^i \text{ for } \bar{i}(t) \geq i > \underline{i}(t) \text{ and} \quad (\text{A2})$$

$$B_{t+1}^{\underline{i}(t)} = R^* B_t^{\underline{i}(t)} - J_t^{\underline{i}(t)} \quad (\text{A3})$$

govern the evolution of the IRA account balances for working and retired members respectively. IRA contributions of working-age individuals are restricted by:

$$0 \leq I_t^i \leq \bar{I}. \quad (\text{A4})$$

That is, contributions cannot be negative and may not exceed a legal limit of $\bar{I}$. A retired individual faces the minimum withdrawal constraint

$$J_t^i \geq F(R^*, \tau i - t) B_t^i, \quad (\text{A5})$$

where $F(R^*, \tau i - t)$ is a fraction of the available IRA balance—and a function of the IRA's rate of return and the individual's life expectancy. In the U.S., the minimum IRA withdrawal rate is the inverse of the individual's life expectancy as determined by the IRS. We approximate this by setting

$$F(R, \kappa) = \frac{R^\kappa (R - 1)}{R^\kappa - 1},$$

which is the current annuity per dollar when the interest rate is R and life expectancy is κ more years.

The balance in the ordinary saving account is denoted A_t^i . Only the standard borrowing constraint $A_t^i \geq 0$ limits its evolution. The budget constraint for each dynasty member is

$$A_{t+1}^i + C_t^i + M_t^i + D_t^i + T_t + I_t^i \leq W_t N^i + J_t^i + R A_t^i. \quad (\text{A6})$$

Here, D_t^i is i 's transfers to other dynasty members, T_t equals the lump-sum tax bill, $N^i = N$ for $i > \underline{i}(t)$, and $N^i = 0$ for $i = \underline{i}(t)$. By definition, $\sum_{i=\underline{i}(t)}^{\bar{i}(t)} D_t^i = 0$.

A.3 Optimal Consumption and Savings

The dyanasty's optimization problem coordinates within-dynasty transfers and savings decisions to maximize total dynastic welfare subject to each individual's budget and asset accumulation constraints and the dynasty's constraint that total within-dynasty transfers equal zero. The necessary conditions for an optimum characterize the allocation across time, members, and goods.

$$C_{t+1}^i \geq \beta R C_t^i, \quad (\text{A7})$$

$$C_t^i = C_t^{i'}, \quad (\text{A8})$$

$$M_t^i = C_t^i \mu_t^i, \quad (\text{A9})$$

for $t = 0, 1, \dots$, and $i, i' = \underline{i}_t, \dots, \bar{i}(t)$. Condition (A7) is the Euler equation allowing for an occasionally binding borrowing constraint. Altruism requires the dynasty to equate consumption across members in condition (A8), and condition (A9) equates the marginal utility of end-of-life medical care with that for any member's ordinary consumption. Incorporating the optimal equalization of consumption across members from (A8) into the utility function in (A1), the dynasty's utility function can be expressed as

$$\sum_{t=0}^{\infty} \beta^t (\ln C_t + \mu_t \ln M_t). \quad (\text{A10})$$

Where $\mu_t = \mu_t^{i(t)}$.

We look for an optimum satisfying the following two conjectures:

Conjecture 4 *Working members' IRA contributions equal their allowed maximum: $I_t^i = \bar{I}$ for $i = \underline{i}_t + 1, \dots, \bar{i}(t)$.*

Conjecture 5 *Retired members' IRA withdrawals equal their required minimum:*

$$J_t^{i(t)} = F(R^*, \tau i - t) B_t^{i(t)}.$$

If Conjecture 4 holds good, then an individual's IRA balance upon retirement equals

$$B^* = \bar{I} \sum_{j=1}^{(\gamma-1)\tau} (R^*)^j.$$

Conjecture 5 requires the withdrawals during retirement to equal $F(R^*, \tau)B^*$. We now define the dynasty's *net earnings* as $Y_t = (\gamma - 1)(W_t N - \bar{I}) + F(R^*, \tau)B^* - \gamma T_t$. With this we can write the dynasty's budget constraint as

$$A_{t+1} + C_t + M_t \leq Y_t + R A_t; \quad (\text{A11})$$

where A_t , C_t and M_t sum the ordinary savings, consumption, and end-of-life care across the dynasty's members.

Note that maximizing (A10) subject to (A11) and $A_t \geq 0$ is identical to the basic model's problem presented in Section 3.1, subject only to the redefinition of net income. Thus, the equivalence of this dynastic model with the infinite-horizon model in the text hinges on whether or not Conjectures 4 and 5 apply: Under these two conjectures, income is exogenous as in the model in the text.

A.3.1 The Deterministic Cycle

We assume here that $W_t = W$ and $T_t = T$. If Conjectures 4 and 5 hold, the present model generates a deterministic cycle in the same way as shown in Section 3.2, with $Y = (\gamma - 1)(W - T - \bar{I}) + \bar{J}$. In particular, Proposition 1 shows that either the borrowing constraint always binds ($\mu \leq \hat{\mu}^\tau$) or that there exists a $\kappa^* < \tau$ such that $A^\kappa > 0$ if and only if $\kappa \geq \kappa^*$ ($\mu > \hat{\mu}^\tau$). For $\kappa < \kappa^*$, $C^\kappa = Y$. When $A^\kappa > 0$, the Euler equation requires consumption to shrink at the rate $C^{\kappa+1}/C^\kappa = \beta R$. We use these results here to derive a condition on R^* , R , and β sufficient for these conjectures to hold in the model's deterministic cycle.

Confirming Conjecture 4 requires to verify that the benefit of the marginal IRA investment in each of the $(\tau - 1)\gamma$ years of an individual's working life exceeds its present utility cost. Similarly, Conjecture 5 requires us to verify that the benefits of leaving an investment in the IRA during retirement exceeds the cost of doing so.

We first examine an individual's IRA contribution in the last working year. If the individual expects to make always the minimum withdrawals from the IRA during retirement, then investing the maximum contribution at the end of one's career squares with the dynasty's utility maximization if and only if

$$\frac{1}{C^\tau} \leq F(R, \tau) \left(\beta \frac{1}{C^1} + \beta^2 \frac{1}{C^2} + \cdots + \beta^\tau \frac{1}{C^\tau} \right). \quad (\text{A12})$$

The left-hand side of (A12) is the utility cost of the marginal IRA contribution, and its right-hand side equals the future utility gain from the contribution given that minimum IRA withdrawal always binds.

We proceed by multiplying both sides of (A12) by C^τ . The Euler inequality in (A7) bounds C^τ/C^κ from below by $(\beta R)^{\tau-\kappa}$, so replacing each consumption ratio on the right-hand side with its corresponding lower bound yields a lower bound for the entire right-hand side. Thus, we conclude that (A12) is true if

$$1 \leq F(R^*, \tau) \beta^\tau (R^{\tau-1} + R^{\tau-2} + \dots + R + 1). \quad (\text{A13})$$

Dividing and multiplying the right hand side by R^τ and solving the resulting sum of a geometric series expresses this inequality as

$$1 \leq (\beta R)^\tau \frac{F(R^*, \tau)}{F(R, \tau)}. \quad (\text{A14})$$

In summary, (A14) guarantees that an individual at the brink of retirement makes the largest allowable contribution to his IRA. It turns out that this condition also guarantees that he makes the maximum contribution in all earlier years and withdraws the results during retirement at the slowest possible rate. We show this in the the following proposition, which ensures, as mentioned above, that there is a nonstochastic cycle in this model.

Proposition 6 *The inequality in (A14) guarantees that Conjectures 4 and 5 hold good.*

Proof. The proof proceeds in three steps. First, we show that (A14) implies that $\beta R^* \geq 1$, i.e., the IRA's rate of return dominates impatience. Second, this guarantees that retirees withdraw from their IRAs as slowly as possible (Conjecture 5). Third, we demonstrate that the ratio of utility benefit to utility cost for any IRA contribution weakly exceeds the analogous ratio from (A12) (Conjecture 4).

The IRA's return dominates impatience ($\beta R^* > 1$). In the noncyclical case of $\tau = 1$, this immediately follows from (A14). Hence, we proceed with $\tau > 1$. Defining

$$g(R, \kappa) \equiv \frac{R - 1}{R^\kappa - 1}, \quad (\text{A15})$$

the inequality of (A14) can be expressed as

$$1 \leq (\beta R^*)^\tau \left(\frac{g(R^*, \tau)}{g(R, \tau)} \right). \quad (\text{A16})$$

The derivative of the function g with respect to R is negative if

$$\kappa \frac{R-1}{R} - \frac{R^\kappa - 1}{R^\kappa} > 0. \quad (\text{A17})$$

The left-hand side of (A17) equals zero when evaluated at $R = 1$, and for $R > 1$, the derivative of the left-hand side of (A17) with respect to R equals

$$\frac{\kappa}{R^2} - \frac{\kappa}{R^{\kappa+1}} > 0. \quad (\text{A18})$$

Hence, for $R > 1$, (A17) holds. This in turn implies that g declines with R , and therefore $g(R^*, \tau) < g(R, \tau)$ given that $R^* > R$. Because the ratio of values of g in (A16) is less than one, $\beta R^* > 1$.

IRA withdrawals equal their minima. We proceed to verify that $\beta R^* > 1$ implies that Conjecture 5 holds. The minimum withdrawal leaves no room for choice in an individual's final year of life, but we must still consider the earlier years of retirement.

Starting from the minimum required, let us consider a marginal increase in the IRA withdrawal. If the borrowing constraint does not bind in that year, then the dynasty is indifferent between the marginal utility of consuming one more unit in the present year at the expense of the utility loss of consuming R less next year (the Euler condition holds with equality). However, the additional withdrawal from the IRA implies a cost of $R^* > R$ next year, and hence the dynasty will not withdraw the additional unit.

If the borrowing constraint does bind, consumption levels in the present and next year both equal Y . Thus, the relevant utility comparison involves only consumption units. The withdrawal produces one additional unit of consumption this year, at the expense of a present value loss of βR^* next year. Since this exceeds one, the dynasty will not withdraw the additional unit in this case either. Hence, Conjecture 5 must hold.

IRA contributions equal their maxima. The discussion preceeding Proposition 6 demonstrates that (A14) guarantees that a worker's final IRA contribution equals its maximum. We proceed to consider the contributions made κ years before retirement for $1 < \kappa \leq \tau$, i.e., in earlier years of the last expenditure cycle before retirement. The benefit of the marginal investment made in this year exceeds its cost if

$$\frac{1}{C^\kappa} \leq (\beta R^*)^{\tau-\kappa} F(R^*, \tau) \left(\beta \frac{1}{C^1} + \beta^2 \frac{1}{C^2} + \cdots + \beta^\tau \frac{1}{C^\tau} \right). \quad (\text{A19})$$

Multiplying both sides through by C^τ and using the Euler inequality in (A7)—as in the discussion of equations (A12)-(A14)—yields

$$\frac{C^\tau}{C^\kappa} \leq (\beta R^*)^{\tau-\kappa} \frac{F(R^*, \tau)}{F(R, \tau)}. \quad (\text{A20})$$

Since $\beta R^* > 1$, the right-hand side of (A19) exceeds the right-hand side of (A14). Furthermore, its left-hand side is weakly less than one because $C^\tau \leq C^\kappa$. Therefore, (A14) implies that (A19) holds good.

The remaining case to consider is a marginal IRA contribution during an earlier expenditure cycle. (That is, when the individual's working life has more than τ years left.) Suppose that it occurs in year κ of that cycle. Its contemporaneous utility cost is the same as that in (A19), while its appropriately discounted utility benefit is that on the right-hand side of (A19) *multiplied by* $(\beta R^*)^{(\tau-\kappa)j}$ for some $j > 1$. Therefore, the condition guaranteeing that the marginal IRA contributions in the final τ years of work increase utility ensures that earlier contributions do as well. Thus, Conjecture 4 is correct if (A14) holds good. ■

Finally, the next proposition shows that (A14) is consistent with the basic assumption of impatience.

Proposition 7 *Given any $R > 1$ and $R^* > R$, there exist values of $\beta < 1/R$ such that (A14) holds good.*

Proof. The proof simply demonstrates the intuitive result that $F(R, \tau)$ is increasing in R , so that if we set βR arbitrarily close to 1 then (A14) holds. Let $f(R, \tau) \equiv \ln(F(R, \tau)) = \kappa \ln R + \ln(R - 1) - \ln(R^\kappa - 1)$. Its first derivative is

$$\begin{aligned} \frac{\partial f(R, \tau)}{\partial R} &= \frac{\kappa}{R} + \frac{1}{R-1} - \frac{\kappa R^{\kappa-1}}{R^\kappa - 1} \\ &= \frac{1}{R-1} - \frac{\kappa}{R} \frac{1}{R^\kappa - 1} \end{aligned}$$

Evaluated at $\kappa = 1$, this derivative equals $1/R$. Therefore, if we can show that $\partial^2 f(R, \kappa)/(\partial R \partial \kappa) > 0$, then we know that the first derivative of interest must be positive. This second derivative equals

$$\frac{\partial^2 f(R, \kappa)}{\partial R \partial \kappa} = \frac{1}{R} \frac{1}{R^\kappa - 1} \left(\kappa \frac{R^\kappa}{R^\kappa - 1} - 1 \right).$$

The term in parentheses is positive because $\kappa R^\kappa > R^\kappa - 1$ for $\kappa > 1$. Therefore, $\partial f(R, \kappa)/\partial R > 0$. Hence, any

$$\beta \in \left[\left(\frac{F(R, \tau)}{R F(R^*, \tau)} \right)^{1/\tau}, \frac{1}{R} \right)$$

satisfies both the initial assumption that $\beta R < 1$ and (A14). Since $F(R^*, \tau) > F(R, \tau)$, this interval is non-empty. ■

If we use $R = 1.04$ from Table 4, set the marginal tax rate at $\lambda = .30$, and assume that $R^* = 1 + (R - 1)/(1 - \lambda)$, then $R^* = 1.0595$. With these rates, the value of β from Table 4 $1/1.055$ satisfies (A14).

A.3.2 The Marginal Propensity to Consume

With the condition ensuring that the dynasty fully exploits its IRA investment opportunities in the deterministic cycle established, we may proceed to consider the dynasty's marginal propensity to consume. When the dynasty *strictly* prefers to exploit fully its IRA investment opportunities, a marginal change to the dynasty's net earnings (either presently or in the future) leaves this choice unchanged. In this case, we can calculate the optimal consumption and savings responses while holding the IRA contributions and withdrawals fixed. As noted above, the dynasty's utility maximization problem is isomorphic to that of the model in the text. Hence, the analysis of small policy changes in this model is the same as in Section 3.

A.3.3 Durable Goods

The quantitative results in the text come from a version of the model with durable goods and collateralized household debt. Extending the dynastic model to include durable goods is straightforward. Each individual accumulates and maintains a stock of durable goods, and the oldest generation bequeaths its stock to the newly-born youngest generation. As with nondurable goods, optimality requires the dynasty to distribute durable goods services equally across its living members. If we impose the maximum IRA contributions and minimum IRA withdrawals on that dynasty, then its utility maximization problem is identical to that considered in Section 3.5. In that case, we can interpret the computational results reported in the text as the outcome of the dynasty's utility maximization problem. It is not hard to show that the inequality in (A14) guarantees that the dynastic analogue to the infinite-horizon model with durable goods has a deterministic cycle in which Conjectures 4 and 5 hold good if nondurable consumption in the infinite horizon model rises only immediately after the periodic expenditure. This is so for the parameter values we use, so the dynastic analysis of this Appendix applies to the model used to create our quantitative results.

A.4 Calibration

Assume that the IRA contribution ceiling is such that the resulting annuity equals labor income net of IRA contributions ($\bar{I}$) of a working member. That is

$$WN - \bar{I} = F(R^*, \tau) \bar{I} \times \sum_{j=0}^{\tau(\gamma-1)} (R^*)^j. \quad (\text{A21})$$

This says that IRA withdrawals are sufficient to fund all regular consumption of retirees plus their proportional share of end-of-life expenditures ($100/\gamma$ percent). Although income and expenditures are equalized, the individual members' assets are still indeterminate, given that different asset distributions can be supported by appropriate transfer schemes. This indeterminacy is resolved by the additional assumption that the only transfers across members are workers' contributions to retirees' medical expenditures. This is realistic for the United States, given that large intervivos transfers not for educational or medical expenses are subject to gift taxes. These motivate avoiding unnecessary transfers. With equal income and equal expenditures year by year, non-retirement assets are also equalized across working and retired members of the dynasty:

$$A_t^{i(t)} = A_t^{i(t)+1} = \dots = A_t^{\bar{i}_t}. \quad (\text{A22})$$

These assumptions have direct bearing on the calibration of the parameters μ and τ . Because retirement saving is fully accounted for by formal retirement saving, non-retirement assets are only for the periodic expenditure. The average balance should then reflect the utility parameter μ as well as the time between expenditures, τ . Given that all variables in the model are proportional to exogenous net labor income, we can calibrate these parameters so that the model's ratio

$$\frac{A_t^i}{WN - T - \bar{I}} = \frac{A_t}{Y}$$

equals the corresponding ratio for working age individuals. Because this model is identical to the model in the text, we adopt this procedure for the calibration in Section 4.2. The actual value of $\bar{I}$ in the United states is probably insufficient to finance all expenditures in retirement, but the dynastic model omits for simplicity pay-as-you-go public pensions. Incorporating a transfer to retirees financed by payroll taxes on younger generations is straightforward and alters no substantial result. In light of this, we find (A21) empirically plausible.

References

- Aizcorbe, A. M., A. B. Kennickell, and K. B. Moore (2003, January). Recent changes in U.S. family finances: Evidence from the 1988 and 2001 Survey of Consumer Finances. *Federal Reserve Bulletin* 89, 1–32.
- Baumol, W. J. (1952). The transactions demand for cash: An inventory theoretic approach. *Quarterly Journal of Economics* 67(4), 545–556.
- Campbell, J. and A. Deaton (1989). Why is consumption so smooth? *The Review of Economic Studies* 56(3), 357–373.
- Campbell, J. R. and Z. Hercowitz (2009). Welfare implications of the transition to high household debt. *Journal of Monetary Economics* 56(1), 1–16.
- Carroll, C. D. (2001). A theory of the consumption function, with and without liquidity constraints. *The Journal of Economic Perspectives* 15(3), 23–45.
- Carroll, C. D., J. Slacalek, and M. Sommer (2009, October). International evidence on sticky consumption growth. Working paper, Johns Hopkins University, the European Central Bank, and the International Monetary Fund.
- De Nardi, M., E. French, and J. B. Jones (2009). Why do the elderly save? The role of medical expenses. Federal Reserve Bank of Chicago Working Paper 2009-02.
- Deaton, A. (1987). Life-cycle models of consumption: Is the evidence consistent with the theory? In T. F. Bewley (Ed.), *Advances in Econometrics: Fifth World Congress, Vol II*, Econometric Society Monographs, pp. 121–148. Cambridge University Press.
- Dynan, K. E. (2000). Habit formation in consumer preferences: Evidence from panel data. *American Economic Review* 90(3), 391–406.
- Flavin, M. (1985). Excess sensitivity of consumption to current income: Liquidity constraints or myopia? *The Canadian Journal of Economics / Revue canadienne d'Economique* 18(1), 117–136.
- Galí, J. (1990). Finite horizons, life-cycle savings, and time-series evidence on consumption. *Journal of Monetary Economics* 26(3), 433–452.

- Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). Household expenditure and the income tax rebates of 2001. *The American Economic Review* 96(5), 1589–1610.
- Lusardi, A. (1996). Permanent income, current income, and consumption: Evidence from two panel data sets. *Journal of Business & Economic Statistics* 14(1), 81–90.
- Parker, J. A. (1999). The reaction of household consumption to predictable changes in social security taxes. *The American Economic Review* 89(4), 959–973.
- Shapiro, M. D. and J. Slemrod (2003). Consumer response to tax rebates. *The American Economic Review* 93(1), 381–396.
- Souleles, N. S. (1999). The response of household consumption to income tax refunds. *The American Economic Review* 89(4), 947–958.
- Tobin, J. (1956). The interest elasticity of the transactions demand for cash. *Review of Economics and Statistics* 38(3), 241–247.
- Zeldes, S. P. (1984). *Optimal Consumption with Stochastic Income*. Ph. D. thesis, Massachusetts Institute of Technology.
- Zeldes, S. P. (1989). Consumption and liquidity constraints: An empirical investigation. *The Journal of Political Economy* 97(2), 305–346.