

NBER WORKING PAPER SERIES

A THEORY OF THE CONSUMPTION FUNCTION,
WITH AND WITHOUT LIQUIDITY CONSTRAINTS
(EXPANDED VERSION)

Christopher D. Carroll

Working Paper 8387
<http://www.nber.org/papers/w8387>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2001

This is a more rigorous and detailed version of a paper written for a Journal of Economic Perspectives symposium on consumption and saving behavior, for publication in the summer of 2001. The JEP version of the paper is intended for a general audience; the version here would be more appropriate for the consumption segment of a first-year graduate course, or more generally as an introduction to modern consumption theory for someone interested in beginning to pursue research in this area. To help new researchers get up to speed, the Mathematica programs that produced all of the theoretical results in this paper are available on the author's website, www.econ.jhu.edu/people/carroll. I am grateful to Carl Christ and to the JEP editors for valuable comments on earlier versions of this paper. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research.

© 2001 by Christopher D. Carroll. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Theory of the Consumption Function, With and Without Liquidity Constraints
(Expanded Version)
Christopher D. Carroll
NBER Working Paper No. 8387
July 2001
JEL No. A23, B22, D91, E21

ABSTRACT

This paper argues that the modern stochastic consumption model, in which impatient consumers face uninsurable labor income risk, matches Milton Friedman's (1957) original description of the Permanent Income Hypothesis much better than the perfect foresight or certainty equivalent models did. The model can explain the high marginal propensity to consume, the high discount rate on future income, and the important role for precautionary behavior that were all part of Friedman's original framework. The paper also explains the relationship of these questions to the Euler equation literature, and argues that the effects of precautionary saving and liquidity constraints are often virtually indistinguishable.

Christopher D. Carroll
Department of Economics
Johns Hopkins University
Baltimore, MD 21218-2685
and NBER
ccarroll@jhu.edu

1 Introduction

Fifteen years ago, Milton Friedman’s 1957 treatise *A Theory of the Consumption Function* seemed badly dated. Dynamic optimization theory not been employed much in economics when Friedman wrote, and utility theory was still comparatively primitive, so his statement of the “permanent income hypothesis” never actually specified a formal mathematical model of behavior derived explicitly from utility maximization. Instead, Friedman relied at crucial points on intuition and verbal descriptions of behavior. Although these descriptions sounded plausible, when other economists subsequently found multiperiod maximizing models that could be solved explicitly, the implications of those models differed sharply from Friedman’s intuitive description of his ‘model.’ Furthermore, empirical tests in the 1970s and ’80s often rejected these rigorous versions of the permanent income hypothesis, in favor of an alternative hypothesis that many households simply spent all of their current income.

Today, with the benefit of a further round of mathematical (and computational) advances, Friedman’s (1957) original analysis looks more prescient than primitive. It turns out that when there is meaningful uncertainty in future labor income, the optimal behavior of moderately impatient consumers is much better described by Friedman’s original statement of the permanent income hypothesis than by the later explicit maximizing versions. Furthermore, in a remarkable irony, much of the empirical evidence that rejected the permanent income hypothesis as specified in tests of the 1970s and ’80s is actually consistent both with Friedman’s original description of the model and with the new version with serious uncertainty.

There are four key differences between the explicit maximizing models developed in the 1960s and ’70s and Friedman’s model as stated in *A Theory of the Consumption Function* (and its important clarification in Friedman (1963)).

First, Friedman repeatedly acknowledged the importance of precautionary saving against future income uncertainty. In contrast, the crucial assumption that allowed subsequent theorists to solve their formal maximizing models was that labor income uncertainty had no effect on consumption, either because uncertainty was assumed not to exist (in the “perfect foresight” model) or because the utility function took a special form that ruled out precautionary motives (the “certainty equivalent” model).¹

Second, Friedman asserted that his conception of the permanent income hypothesis implied that the marginal propensity to consume out of transitory “windfall” shocks to income was about a third. However, the perfect foresight and certainty equivalent models typically implied an MPC of 5 percent or less.

Third, Friedman (1957) asserted that the “permanent income” that determined

¹The uncertainty considered here is explicitly labor income uncertainty. Samuelson (1969) and Merton (1969) found explicit solutions long ago in the case where there is rate-of-return uncertainty but no labor income uncertainty, and showed that rate-of-return uncertainty does not change behavior much compared to the perfect-foresight model.

current spending was something like a mean of the expected level of income in the very near-term: “It would be tempting to interpret the permanent component [of income] as corresponding to the average lifetime value . . . It would, however, be a serious mistake to accept such an interpretation.” He goes on to say that households in practice adopt a much shorter ‘horizon’ than the remainder of their lifetimes, as captured in the assumption in Friedman (1963) that people discount future income at a “subjective discount rate” of 33-1/3 percent. In contrast, the perfect foresight and certainty equivalent models assumed that future income was discounted to the present at market interest rates (say, 4 percent).

Finally, as an interaction between all of the preceding points, Friedman indicated that the *reason* distant future labor income had little influence on current consumption was “capital market imperfections,” which encompassed both the fact that future labor income was uninsurably uncertain and the difficulty of borrowing against such income (for example, see Friedman (1963) p. 10).

It may seem remarkable that simply adding labor income uncertainty can transform the perfect foresight model into something closely resembling Friedman’s original framework; in fact, one additional element is required to make the new model generate Friedmanesque behavior: Consumers must be at least moderately impatient. The key insight is that the precautionary saving motive intensifies as wealth declines, because poorer consumers are less able to buffer their consumption against bad shocks. At some point, the intensifying precautionary motive becomes strong enough to check the decline in wealth that would otherwise be caused by impatience. The level of wealth where the tug-of-war between impatience and prudence reaches a stalemate defines a ‘target’ for the buffer stock of precautionary wealth, and many of the insights from the new model can best be understood by considering the implications and properties of this target.

A final insight from the new analysis is that precautionary saving behavior and liquidity constraints are intimately connected.² Indeed, for many purposes the behavior of constrained consumers is virtually indistinguishable from the behavior of unconstrained consumers with a precautionary motive; average behavior depends mainly on the degree of impatience, not on the presence or absence of constraints. As a result, most of the existing empirical studies that supposedly test for constraints should probably be reinterpreted as evidence on the average degree of impatience. Furthermore, future studies should probably focus more directly on attempting to measure the average degree of impatience rather than on attempting to detect constraints.

²For a rigorous analysis of the relationship between constraints and precautionary behavior, see Carroll and Kimball (2001).

2 The Modern Model(s) of Consumption

Current graduate students rarely appreciate how difficult it was to forge today's canonical model of consumption based on multiperiod utility maximization. The difficulty of the enterprise is attested by the volume of literature devoted to the problem from the 1950s through the '70s, beginning with the seminal contribution of Modigliani and Brumberg (1954). The model that eventually emerged has several key characteristics. Utility is time separable; that is, the utility that consumption yields today does not depend on the levels of consumption in other periods, past or future. Future utility is discounted geometrically, so that utility one period away is worth β units of this period's utility, utility two periods away is worth β^2 , and so on, for some β between 0 and 1. Furthermore, the utility function must satisfy various criteria of plausibility like decreasing marginal utility, decreasing absolute risk aversion, and so on. Finally, the model must incorporate a mathematically rigorous description of how noncapital income, capital income, and wealth evolve over time.

A version of the maximization problem inherited from this literature can be written as follows. A consumer in period t (who has already been paid for period t 's labor) has an amount of total resources X_t ('cash-on-hand' in Deaton's (1991) terminology), the sum of this period's wealth and this period's labor income. Given this starting position, the consumer's goal is to maximize expected discounted utility from consumption between the current period t and a final period of life T ,

$$\max E_t \left[\sum_{s=t}^T \beta^{s-t} u(\tilde{C}_s) \right] \quad (1)$$

(where the \sim over C_s indicates that its value may be uncertain as of the date at which expectations are being taken) subject to a set of budget constraints and shocks,

$$W_{s+1} = R_{s+1}(X_s - C_s) \quad (2)$$

$$Y_{s+1} = P_{s+1}\epsilon_{s+1} \quad (3)$$

$$P_{s+1} = GP_s N_{s+1} \quad (4)$$

$$X_{s+1} = W_{s+1} + Y_{s+1} \quad (5)$$

where beginning-of-period wealth next period, W_{t+1} , is equal to unspent resources from period t accumulated at a (potentially uncertain) gross interest rate R_{t+1} ; Y_{t+1} is labor (or more properly 'noncapital') income in period $t + 1$, which is equal to 'permanent labor income' P_{t+1} multiplied by a mean-one transitory shock ϵ_{t+1} , $E_t[\tilde{\epsilon}_{t+1}] = 1$; permanent labor income grows by a factor G between periods and is also potentially subject to shocks, N_{s+1} ; and 'cash-on-hand' in period $t + 1$ is equal to beginning-of-period wealth W_{t+1} plus the period's labor income Y_{t+1} .

One of the unpleasant discoveries in the 1960s and '70s was that when there is uncertainty about the future level of labor income (i.e. if ϵ and N have variances

greater than zero), it appears to be impossible (under plausible assumptions about the utility function, e.g. constant relative risk aversion $u(c) = c^{1-\rho}/(1-\rho)$) to derive an explicit solution for consumption as a direct (analytical) function of the model's parameters. This is not to say that nothing at all is known about the structure of optimal behavior under uncertainty; for example, it can be proven that consumption always rises in response to a pure increment to wealth. But an explicit solution for consumption is not available.

2.1 The Perfect Foresight/Certainty Equivalent Model

Economists' main response to this problem was to focus on two special cases where the model can be solved analytically: The "perfect foresight" version in which uncertainty is simply assumed away, or the "certainty equivalent" version in which consumers are assumed to have quadratic utility functions (despite unattractive implications of quadratic utility like risk aversion that increases as wealth rises, and the existence of a 'bliss point' beyond which extra consumption reduces utility).

The perfect foresight and certainty equivalent solutions are very similar; for brevity, I will summarize only the perfect foresight solution, in which the optimal level of consumption is directly proportional to total "wealth," which is the sum of market wealth W_t and 'human wealth' H_t ,

$$C_t = k_t(W_t + H_t), \tag{6}$$

where market wealth W_t is real and financial capital, while human wealth is mainly current and discounted future labor income (though in principle H_t also includes the discounted value of transfers and any other income not contingent on saving decisions; henceforth I refer to these collectively as 'noncapital' income). The constant of proportionality, k_t , depends the time preference rate, the interest rate, and other factors.

A simple example occurs when consumers care exactly as much about future utility as about current utility ($\beta = 1$); the interest rate is zero; and there is no current or future noncapital income ($H_t = 0$). In this case, the optimal plan is to divide existing wealth evenly among the remaining periods of life. If we assume an average age of death of 85, this model implies that the marginal propensity to consume out of shocks to wealth for consumers younger than 65 should be less than $(1/20)$, or 5 percent – since the change in wealth will be spread evenly over at least 20 years. Furthermore, the theory implies that the MPC out of unexpected transitory shocks to noncapital income ("windfalls"; e.g. finding a \$100 bill in the street) is the same as the MPC out of wealth, because once the windfall has been received, it is theoretically indistinguishable from the wealth the consumer already owned. When the model is made more realistic by allowing for positive interest rates, consumers younger than 65, etcetera, it still implies that the average MPC should be quite low, generally less than 0.05.

In contrast, Friedman (1963) asserted that his conception of the permanent income hypothesis implied an MPC out of transitory shocks of about 0.33 for the typical consumer.³ Friedman (1963) provided an extensive summary of the existing empirical evidence tending to support the proposition of an MPC of roughly a third. From today's perspective, however, the most surprising aspect of Friedman's (1957, 1963) arguments is that their main thrust is to prove an MPC much less than one (to discredit the 'Keynesian' model that said consumption was roughly equal to current income), rather than to prove an MPC significantly greater than 0.05.

The 15 years after the publication of *A Theory of the Consumption Function* produced many studies of the MPC. Particularly interesting were some natural experiments. In 1950, unanticipated payments were made to a subset of U.S. veterans holding National Service Life Insurance policies; the marginal propensity to consume out of these dividends seems to have been between about 0.3 and 0.5. Another natural experiment was the reparations payments certain Israelis received from Germany in 1957-58.⁴ The marginal propensity to consume out of these payments appears to have been around 20 percent, with the lower figure perhaps accounted for by the fact that the reparations payments were very large (typically about a year's worth of income).⁵ On the whole, these studies were viewed at the time as supporting Friedman's model because the estimated MPCs were much less than one.

The change in the profession's conception of the permanent income hypothesis in the 1970s from Friedman's (1957, 1963) version to the perfect foresight/certainty equivalent versions (with their predictions of an MPC of 0.05 or less) is nicely illustrated by a well-known paper by Hall and Mishkin (1982) that found evidence of an MPC of about 0.2 using data from the *Panel Study of Income Dynamics* (PSID). Rather than treating this as evidence in favor of a Friedmanesque interpretation of the permanent income hypothesis, the authors concluded that at least 15-20 percent of consumers *failed* to obey the PIH because their MPCs were much greater than 0.05.

³My definitions of 'transitory' and 'permanent' shocks (spelled out explicitly in the next subsection) correspond to usage in much of the modern consumption literature, but differ from Friedman's (1957) usage. In fact, Friedman (1957) actually states that the MPC out of 'transitory income shocks' is zero, but Friedman (1963) was very clear that in his conception of the PIH, first-year consumption out of windfalls was about 0.33. The reconciliation is that such windfalls were not 'transitory' shocks in Friedman's terminology. Terminology aside, Friedman's quantitative predictions for how consumption should change, for example in response to a windfall, are clear, so I will simply translate the Friedman model's predictions into modern terminology without further remark, e.g. by stating that Friedman's model implies that the MPC out of (my definition of) transitory shocks is a third.

⁴For an excellent summary of these studies by Bodkin (1959), Kreinin (1961), Landsberger (1966), and others see Mayer (1972).

⁵The concavity of the consumption function discussed below, and proved in Carroll and Kimball (1996), implies that the MPC out of a large shock should be smaller than the MPC out of a small shock.

2.2 The New Model

The principal development in consumption theory in the last 15 years or so, starting with Zeldes (1984), is that spectacular advances in computer speed have allowed economists to relax the perfect foresight/certainty equivalence assumption and determine optimal behavior under realistic assumptions about uncertainty.

A preliminary step was to determine the characteristics of the income uncertainty that typical households face.⁶ Using annual income data for working-age households participating in the PSID, Carroll (1992) found that the household noncapital income process is well approximated as follows. In period t a household has a certain level of ‘permanent noncapital income’ P_t , which is defined as the level of noncapital income the household would have gotten in the absence of any transitory shocks to income.⁷ Actual income is equal to permanent income multiplied by a transitory shock, $Y_t = P_t\epsilon_t$ where permanent income P_t grows by a factor G over time, $P_t = GP_{t-1}$. Each year there is a small chance (probability 0.005) that actual household income will be essentially zero ($\epsilon_t = 0$), typically corresponding in the empirical data to a spell of unemployment or temporary illness or disability. If the transitory shock does not reduce income all the way to zero, that shock is distributed lognormally with a mean value of one and a standard deviation of $\sigma_\epsilon = 0.1$. Carroll (1992) and subsequent papers also find strong evidence for permanent as well as transitory shocks to income, also with an annual standard deviation of perhaps 0.1. However, because permanent shocks complicate the exposition without yielding much conceptual payoff, I will suppress them for the purposes of this paper and compensate by boosting the variance of the transitory component to $\sigma_\epsilon = 0.2$; for the version with both transitory and permanent shocks, see Carroll (1992). The PSID also shows the annual household income growth factor to be about $G = 1.03$ or 3 percent growth per year for households whose head is in the prime earning years of 25-50.

The next step in solving the model computationally is to choose values for the parameters that characterize consumers’ tastes. For the simulation results presented in this paper, I will assume a rather modest precautionary saving motive by choosing a coefficient of relative risk aversion of $\rho = 2$, toward the low end of the range from 1 to 5 generally considered plausible.⁸ I follow a traditional calibration in the macro

⁶One might suppose that this would have been a subject of preexisting research in the labor economics literature. However, labor economists tend to focus on the wage process for individual workers rather than the degree of uncertainty in post-transfer, household-level noncapital income that is the relevant concept from consumption theory.

⁷Friedman (1963), p. 5, says that ‘permanent income is of the nature of the mean of a hypothetical probability distribution’ which is precisely what P_t here is.

⁸This choice of ρ implies that a consumer would be indifferent between consuming \$66,666 with certainty or consuming \$50,000 with probability 1/2 and \$100,000 with probability 1/2. For $\rho = 0$, the consumer is not risk averse at all and would be indifferent between \$75,000 with certainty and \$50,000 with probability .5 and \$100,000 with probability .5. For $\rho = \infty$, the consumer is infinitely risk averse, and would choose \$50,000.01 with certainty over equal probabilities of \$50,000 and \$100,000.

literature and choose a time preference factor of $\beta = 0.96$ implying that consumers discount future utility at a rate of about 4 percent annually, and I make a symmetric assumption that the interest rate is also 4 percent per year.

We are now in position to describe how the model can be solved computationally. As is usual in this literature, it is necessary to solve backwards from the last period of life. For simplicity, we will assume that the income process described above, with constant income growth G , holds for every year of life up to the last. (For a version with a more realistic treatment of the lifetime income profile, including the drop in income at retirement, see Carroll (1997)).

In the last time period, the solution is easy: The benchmark model assumes there is no bequest motive, so the consumer spends everything. Following Deaton (1991), define cash-on-hand X as the sum of noncapital income and beginning-of-period wealth (including any interest income earned on last period's savings). In the second-to-last period of life, the consumer's goal is to maximize the sum of utility from consumption in period $T - 1$ and the mathematical *expectation* of utility from consumption in period T , taking into account the uncertainty that results from the possible shocks to future income Y_T . For any specific numerical levels of cash-on-hand and permanent income in period $T - 1$ (say, $X_{T-1} = 5$ and $P_{T-1} = 1.4$), a computer can calculate the sum of current and expected future utility generated by any particular consumption choice. The optimal level of consumption for $\{X_{T-1}, P_{T-1}\} = \{5, 1.4\}$ can thus be found by a computational algorithm that essentially tries out different guesses for C_{T-1} and homes in on the choice that yields the highest current and discounted expected future utility.

Note that for each different combination of $\{X_{T-1}, P_{T-1}\}$, the utility consequences of many possible choices of C_{T-1} must be compared to find the optimum, and for each C_{T-1} that is *considered*, the numerical expectation of next period's utility must be computed. The solution procedure is basically to calculate optimal C_{T-1} for a grid of many possible $\{X_{T-1}, P_{T-1}\}$ choices, and then to construct an approximate consumption function by interpolation ('connect-the-dots').

Once the approximate consumption rule has been constructed for period $T - 1$, the same steps can be repeated to construct a consumption rule for $T - 2$ and so on.

This begins to give the flavor for why numerical solutions are so computation intensive. Indeed, the problem as just described would be something of a challenge even for current technology. Fortunately, there is a trick that makes the problem an order of magnitude easier: Everything can be divided by the level of permanent income. That is, defining the cash-on-hand *ratio* as $x_t = X_t/P_t$ and $c_t = C_t/P_t$, it is possible to find the optimal value of the consumption-to-permanent-income ratio as a function of the cash-on-hand ratio, so that rather than solving the problem for a two-dimensional grid of $\{X_{T-1}, P_{T-1}\}$ points one can solve for a one-dimensional vector of values of $\{x_{T-1}\}$.




Figure 1: Convergence of Consumption Functions $c_{T-n}(x)$ as n Rises

Formally, the problem can be rewritten in the recursive value function form⁹

$$v_t(x_t) = \max_{\{c_t\}} u(c_t) + \beta E_t[G^{1-\rho} v_{t+1}(\tilde{x}_{t+1})] \quad (7)$$

s.t.

$$w_{t+1} = (R/G)(x_t - c_t) \quad (8)$$

$$x_{t+1} = w_{t+1} + \epsilon_{t+1}. \quad (9)$$

The solution to the optimal consumption problem is depicted in Figure 1. The cash-on-hand ratio x is on the horizontal axis. The optimal consumption ratio for a given cash-on-hand ratio is on the vertical axis. The solid lines represent the consumption rules for different time periods, showing how optimal consumption changes as the ratio of cash-on-hand to labor income increases.

Consumption in the last period $c_T(x)$ coincides with the 45 degree line, indicating consumption equal to cash-on-hand. For very low levels of x , consumption in the second-to-last period $c_{T-1}(x)$ is fairly close to the 45 degree line; the consumer spends almost, but not quite, everything. This reflects the precautionary motive: Because there is a chance the consumer will receive zero income in period T , she will never

⁹See Carroll (2000a) for a proof, and for a detailed description of several other tricks that make the problem computationally tractable.

spend all of her period- $T - 1$ resources because of the dire consequences of arriving at T with nothing and then possibly receiving zero income. Note the contrast with behavior at high levels of wealth; for example, at an x_{T-1} of around 10 the figure shows c_{T-1} of a bit more than 5 – indicating that at this large level of wealth the consumer divides remaining lifetime resources roughly evenly between the last two periods of life.

An important feature of this problem is that, if certain conditions hold (in particular, if consumers are ‘impatient’ in a sense to be described shortly), the successive consumption rules $c_T(x), c_{T-1}(x), c_{T-2}(x), \dots, c_{T-n}(x)$ will ‘converge’ as n grows large. The meaning of convergence is most easily grasped visually: In Figure 1, the rules $c_T(x)$ and $c_{T-1}(x)$ are very far apart, while the rules $c_{T-10}(x)$ and the converged consumption rule $c(x)$ (which can be thought of as $c_{T-\infty}(x)$) are very close.

The importance of convergence can best be understood by contrasting it with the alternative. Modigliani (1966) points out that in the certainty equivalent model, optimal behavior is different at every different age, so that one cannot draw many general lessons about consumption behavior from the rule for any particular age. In the model solved here, however, behavior is essentially identical for all consumers more than 10 years from the end of life, so analysis of the converged consumption rule yields insights about behavior of most agents in the economy.

What is required to generate convergence? Deaton (1991) and Carroll (2001b) show that the necessary condition is that consumers be impatient, in the sense that if there were no uncertainty or liquidity constraints the consumer would choose to spend more than her current income. Technically, the required condition is

$$(R\beta)^{1/\rho} < G, \tag{10}$$

where ρ is the coefficient of relative risk aversion and G is the income growth factor.

Consider the version of this equation where $G = \rho = 1$, so that consumers are impatient if $R\beta < 1$. In this case, impatience depends directly on the whether the reward to waiting, as determined by the interest rate factor R , is large enough to overcome the utility cost to waiting, β . Positive income growth ($G > 1$) makes consumers more impatient (in the sense of wanting to spend more than current income) because forward-looking consumers with positive income growth will want to spend some of their higher future income today. Finally, the exponent $(1/\rho)$ on the $R\beta$ term captures the ‘intertemporal elasticity of substitution,’ which measures the extent to which the consumer responds to the net incentives for reallocating consumption between periods.

The remainder of the paper will focus almost exclusively on implications of the converged consumption function. It is natural to wonder, however, whether we should expect these results to be useful in understanding the behavior of consumers whose permanent income paths over the lifetime do not resemble the “constant growth at rate G until death” specification used here. For instance, income can be predicted to decline at retirement! However, Carroll (1997) shows that when a model like this is solved with an empirically realistic pattern of income growth over the lifetime, the

consumption function resembles the ‘converged’ consumption function examined here until roughly age 50. After 50, with retirement looming, the consumer begins saving substantial amounts and behavior begins more and more to resemble that in the perfect foresight model. Thus, the results in the remainder of the paper based on the converged consumption function are most appropriately represented as characterizing the behavior of moderately impatient households up to about age 50.¹⁰

At present, three further observations about the converged consumption function depicted in Figure 1 are important. (The general shape of the consumption function, and the validity of the points made here, are robust to alternative assumptions about parameter values, so long as consumers remain moderately impatient.)

First, the converged consumption function is everywhere well below the perfect foresight solution (the dashed line). Since precautionary saving is defined as the amount by which consumption falls as a consequence of uncertainty, the difference between the converged $c(x)$ and the dashed perfect-foresight line measures the extent of precautionary saving. The precautionary effect is large here because under our baseline parameter values, human wealth is quite large and therefore induces a lot of consumption by the perfect-foresight consumers. In contrast, consumers with a precautionary motive are unwilling to spend much on the basis of uncertain future labor income, so the large value of human wealth has little effect on their current consumption.

The second important observation is that as x gets large, the slope of $c(x)$ (which is to say, the marginal propensity to consume) gets closer and closer to the slope of the dashed perfect foresight line. That is, as wealth approaches infinity the marginal propensity to consume approaches the perfect foresight MPC. This happens because as wealth approaches infinity the proportion of future consumption that will be financed out of uncertain labor income approaches zero, so the labor income uncertainty becomes irrelevant to the consumption decision.

The final observation is that for periods before the last one the consumption function lies everywhere below the 45-degree line; that is, consumers choose never to borrow (which they would need to do in order to have $c > x$ and to be above the 45-degree line), even though no liquidity constraint was imposed in solving the problem.

This last result deserves explanation. As noted above, in the second-to-last period, consumers will always choose to spend less than their cash-on-hand because of the risk of zero income in the last period of life. If we know that in period $T - 1$ consumption will be less than x , then that implies that in period $T - 2$ the consumer will always behave in such a way to make sure that he arrives in $T - 1$ with positive assets, again out of the fear of a zero-income event in $T - 1$. Similar logic goes through recursively to any earlier period.

This mechanism for preventing borrowing may seem rather implausible, relying as it does on the slight possibility of disastrous zero-income events. However, essentially

¹⁰Recent work by Gourinchas and Parker (1999) finds the switchpoint to be between 40 and 45 rather than 50, but Cagetti’s (1999) similar work suggests a later switching age.

the same logic works as long as income has a well-defined lower bound. For example, suppose the worst possible outcome were that income might fall to, say, 30 percent of its permanent level. In this case the recursive logic outlined above would not prohibit borrowing. But it *would* prevent the consumer from borrowing more than the amount \underline{H} that could be repaid with certainty out of the lowest possible future income stream. In this case, consumers would define their precautionary target in terms of the size of their wealth holdings *in excess of* the lowest feasible level $-\underline{H}$. The distinctive features of the model discussed below would all go through, with the solitary difference that the average level of wealth would be lower (perhaps even negative).

This logic provides the simplest intuition for a fundamental conclusion: The precautionary saving motive can generate behavior that is virtually indistinguishable from that generated by a liquidity constraint,¹¹ because the precautionary saving motive essentially induces self-imposed reluctance to borrow (or borrow too much).

3 Implications

3.1 Concavity of the Consumption Function and Buffer Stock Saving

Perhaps the most striking feature of the converged consumption function $c(x)$ depicted in figure 1 is that the marginal propensity to consume (the slope of the consumption function) is much greater at low levels of cash-on-hand than at high levels. In other words, the converged consumption function is strongly concave.¹² Thus, the first intuitive result that comes out of the analysis is that, as Keynes (1935) argued long ago, rich people spend a smaller proportion of any transitory shock to their income than do poor people.

Carroll (2001b) shows that concavity of the consumption function also implies that impatient consumers will engage in ‘buffer-stock’ saving behavior. That is, there will be some target level of the cash-on-hand ratio x^* such that, if actual cash-on-hand is greater than the target, impatience will outweigh prudence and wealth will fall (formally, $E_t[x_{t+1} < x_t | x_t > x^*]$), while if cash-on-hand is below the target the precautionary saving motive will outweigh impatience and the consumer will try to build wealth up back toward the target (formally, $E_t[x_{t+1} > x_t | x_t < x^*]$). As usual, this result is something that Friedman grasped intuitively: He refers repeatedly to the role of wealth as an ‘emergency reserve’ against uncertainty or a ‘balancing resource’; indeed,

¹¹In fact, Carroll and Kimball (2001) show that as the probability of the zero-income events approaches zero, behavior in the model with zero-income events becomes mathematically identical to behavior in the liquidity-constrained model.

¹²Carroll and Kimball (1996) provide a proof that uncertainty induces a concave consumption function for a very broad class of utility functions, including the constant relative risk aversion form used here.

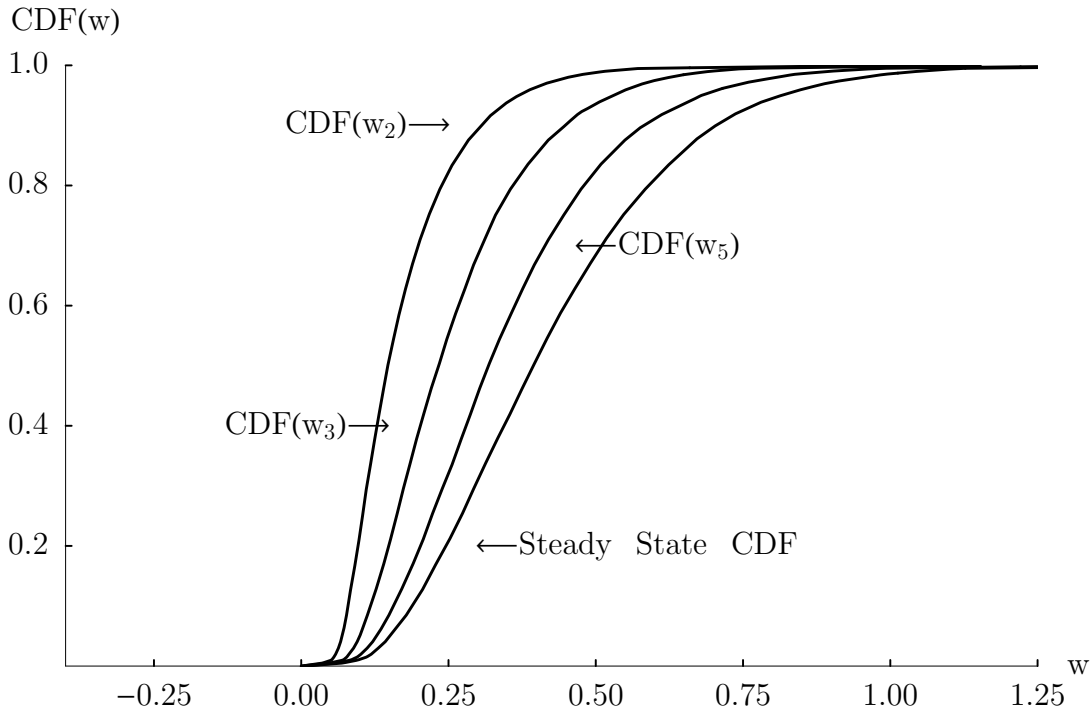


Figure 2: Cumulative Distribution Functions Starting With $w_{1,i} = 0 \forall i$

Mayer (1972), p. 70 summarizes Friedman’s version of the PIH succinctly: ‘It is basic to [Friedman’s] permanent income theory that households attempt to maximize utility by using savings as a buffer against income fluctuations.’

Buffer-stock saving behavior is a qualitative implication of the model. In order to determine the model’s quantitative implications (for example, what it predicts about the average value of the MPC), it is necessary to simulate a population of consumers behaving according to the converged consumption rule. Figure 2 presents the results when a population of 10,000 consumers is endowed with initial wealth $w_{1,i} = 0 \forall i$, then appropriately-distributed random income shocks are drawn to generate $x_{1,i}$, implying consumption $c_{1,i} = c(x_{1,i})$ and second period wealth $w_{2,i} = (R/G)(x_{1,i} - c_{1,i})$, and so on. The figure shows the evolution of the distribution of wealth $w_{t,i}$ in years 2, 3, and 5, along with the steady-state distribution that emerges after sufficiently many periods. It clearly does not take long for the actual wealth distribution to get fairly close to the steady-state distribution, so statistics for consumers distributed according to the steady-state distribution should be a good approximation to typical behavior most of the time (even if the economy is for some reason temporarily out of steady-state).

The first row of Panel A of Table 1 provides a variety of statistics about average behavior when consumers are distributed according to the steady-state distribution generated by the baseline parametric assumptions. Columns two and three indicate that the mean and median of the wealth ratio are both about 0.4, or equal to about five

months' worth of permanent noncapital income (remember that the time unit is a year). The average marginal propensity to consume is 0.33, in the ballpark of both empirical estimates and Friedman's (1957) statement of his conception of the permanent income hypothesis, but a long way from the approximately 0.04 implied by the perfect foresight model under our baseline parameter values.

The second row of Panel A presents results under the assumption that household noncapital income growth is 2 percent a year, rather than the baseline of 3 percent. Lower income growth makes people more "patient," in the sense that the contrast between tomorrow's and today's income – and thus the temptation to borrow against future income – is not as great. The table shows that greater patience leads to a higher mean wealth ratio a lower average MPC.

The final row of Panel A presents results when predictable income growth is zero.¹³ With these extremely patient consumers, who cannot rely on future income gains at all, average wealth is much higher, and the average MPC is only about 0.06, not much greater than in the perfect foresight model.

These results confirm that if consumers are moderately impatient, their behavior in the modern model with uncertainty resembles Friedman's conception of the permanent income hypothesis. Neither liquidity constraints nor myopia is necessary to generate the high average marginal propensity to consume that has repeatedly been found in empirical studies and that Friedman (1957) deemed consistent with his conception of the permanent income hypothesis. Impatience plus uncertainty will do the trick.

The reason precautionary saving increases the MPC is because the precautionary motive relaxes as the level of wealth rises. To put it another way, an extra unit of cash-on-hand today means that one has a better ability to buffer consumption against income shocks in the future, and so there is less need to depress consumption to build up one's precautionary assets. Thus, the decline in the intensity of the precautionary motive as cash-on-hand rises allows consumption to rise faster than it would in the absence of a precautionary motive – which is to say, the MPC out of cash-on-hand (and therefore the MPC out of transitory shocks to income) is higher.

Recall that another difference between Friedman and the subsequent models was in the rate at which consumers were assumed to discount future income. In the subsequent models, the mean expectation of future labor income was discounted to the present at a market interest rate (say, 4 percent). Friedman (1963) insisted that future labor income was discounted at a rate of around 33 percent. (A substantial body of empirical evidence confirms that the actual reaction of consumption to information about future income is much smaller than the perfect foresight and certainty equivalent models imply; see Campbell and Deaton (1989); Viard (1993); Carroll (1994); and the large literature that finds that saving responds much less than one-for-one to expected future pension benefits (Samwick (1995)).)

¹³In this case the consumer is on the edge of failing the impatience condition (but the condition does hold because $(R\beta)^{1/\rho} = 0.9992 < 1.00$ under the baseline values for $\{R, \beta, \rho\} = \{1.04, 0.96, 2\}$).

Table 1: Steady-State Statistics For Alternative Consumption Models

Income Growth Factor	Mean w	Median w	Aggregate Consumption Growth	Mean MPC	Frac With $w < 0$	Frac With $w = 0$
Panel A. Baseline Model, No Constraints						
G=1.03	0.43	0.40	1.030	0.330	0.000	0.000
G=1.02	0.52	0.48	1.020	0.276	0.000	0.000
G=1.00	2.26	2.06	1.000	0.064	0.000	0.000
Panel B. Strict Liquidity Constraints						
G=1.03	0.28	0.24	1.030	0.361	0.000	0.070
G=1.02	0.36	0.32	1.020	0.301	0.000	0.051
G=1.00	2.28	2.06	1.000	0.065	0.000	0.000
Panel C. Borrowing Up To 0.3 Allowed						
G=1.03	-0.03	-0.06	1.030	0.361	0.611	0.000
G=1.02	0.06	0.01	1.020	0.299	0.478	0.000
G=1.00	1.94	1.71	1.000	0.064	0.023	0.000
Panel D. Borrowing Up to 0.3 at $R = 1.15$ Allowed						
G=1.03	0.11	0.07	1.030	0.327	0.320	0.058
G=1.02	0.21	0.16	1.020	0.274	0.210	0.046
G=1.00	2.11	1.89	1.000	0.064	0.007	0.002
Panel E. Statistics from the 1995 SCF						
-	1.02	0.29	-	-	0.205	0.025

Notes: Results in Panels A through D reflect calculations by the author using simulation programs available at the author's website, <http://www.econ.jhu.edu/people/carroll/ccarroll.html>. In Panel A, no constraint is imposed, but income can fall to zero, which prevents consumers from borrowing. In Panels B through D, the worst possible event is for income to fall to half of permanent income. For comparison, Panel E presents the mean and median values of the ratio of nonhousing wealth to permanent income from the 1995 *Survey of Consumer Finances* for non-self-employed households whose head was aged 25-50; the measure of permanent income is actual measured household income for households who reported that their income over the past year was 'about normal', and whose reported income was at least \$5000; other households are dropped. The program that generates these statistics (and figure 6) is also available at the author's website.

We can examine this controversy in the new model by determining how average consumption changes when expectations about the future path of income change. Suppose we have a population of consumers who have received their period t income and are distributed according to the steady-state distribution of x_t that obtains under the baseline parameter values. Now consider informing these consumers that henceforth growth will be $G = 1.02$ rather than 1.03. It turns out that under the baseline parameter values, consumers react to the news of the change in income growth as though they are discounting future noncapital income at a 39 percent rate - even higher than Friedman's estimate of 33 percent!¹⁴ The reason for the high discount rate is that prudent consumers know it would be unwise to spend today on the basis of future income that might not actually materialize.

3.2 The Consumption Euler Equation

Robert Hall (1978) provided the impetus for a large empirical literature over the past two decades by pointing out that in the certainty equivalent model, the predictable change in consumption in a given period should be unrelated to any information that the consumer possessed in earlier periods; consumption should follow a 'random walk.'

To derive this result, Hall relied on an optimality condition known as the Euler equation which links marginal utility in adjacent periods. In the CRRA-utility model with uncertainty, a crude ('first-order') approximation to the Euler equation implies that an equation of the form

$$E_t[\Delta \log C_{t+1}] \approx \rho^{-1}(r - \theta) \quad (11)$$

$$\Delta \log C_{t+1} \approx \rho^{-1}(r - \theta) + \zeta_{t+1} \quad (12)$$

will hold, where $\beta = 1/(1 + \theta)$ and θ is the pure rate of time preference, and ζ_{t+1} is an 'expectational error,' which implies that nothing known in period t should be able to predict the value of ζ_{t+1} .

A more precise ('second order') approximation of the consumption Euler equation

¹⁴The procedure for calculating an average 'effective' interest rate is as follows. First, determine what aggregate consumption would be in period t if consumers continued to expect $G = 1.03$; call the result C_t^{03} . Next, find the converged consumption rule under the expectation that $G = 1.02$, and use it to determine how much consumption would be done if consumers' expectations were suddenly switched to $G = 1.02$ permanently; call that result C_t^{02} . Finally, find the value of the interest factor R such that, in the perfect foresight model, if growth expectations changed from $G = 1.03$ to $G = 1.02$ then consumption would change by $C_t^{03} - C_t^{02}$. Unfortunately, the answer that one gets from this methodology for the "effective" interest rate depends very much on how the change in income is distributed over time, its stochastic properties, the level of current wealth, and all of the other parameters of the model.

leads to a relationship of the form:¹⁵

$$E_t[\Delta \log C_{t+1}] \approx \rho^{-1}(r - \theta) + \left(\frac{\rho + 1}{2}\right) E_t[(\Delta \log C_{t+1})^2]. \quad (13)$$

The term involving the expectation of the square of consumption growth is vanishingly small when there is no uncertainty, so in this case the equation essentially collapses to (11). However, when there is important uncertainty the expected square of consumption growth need not be negligible at all. This term, which resembles a variance, reflects the effect of precautionary saving on consumption growth.

One of the most surprising features of equations (11) and (13) is that the growth rate of income does not appear in either equation. Thus, these equations appear to imply that consumption growth is determined entirely by consumers' tastes and does not depend at all on income growth.

However, Panel A of Table 1 shows that when the growth rate of permanent income is changed from 3 percent to 2 percent to 0 percent, the growth rate of aggregate consumption changes in an identical way, from 3 to 2 to 0 percent. At a minimum, this tells us that there is something profoundly wrong with at least (11) as a way to describe the relationship between income growth and consumption growth.

It turns out that equation (13) is not as hopeless, because it contains a term involving the square of consumption growth. It is clearly possible for expected consumption growth to equal expected income growth for some possible value of the precautionary term $E_t[(\Delta \log C_{t+1})^2]$; in fact, it turns out that this precautionary term is precisely the thing that adjusts to make aggregate consumption growth match aggregate income growth.

The magnitude of the precautionary term for any given x_t can only be determined by solving the model numerically and then computing the expectation numerically. Figure 3 plots the expectation of consumption growth as a function of the level of the period- t cash-on-hand ratio x_t . The most striking thing about the figure is the strong negative relationship between the level of x_t and expected consumption growth. This is a manifestation of the weakening of the precautionary motive as wealth rises. For example, at very low levels of cash-on-hand x_t (levels below x^*), the intense precautionary motive induces the consumer to keep c_t low compared to mean expected future income, out of the fear of an unfavorable income shock in period $t + 1$. But by definition, the actual draw of income in period $t + 1$ is usually not unfavorable, so most of the time the consumer's high precautionary saving in period t will result in a larger x_{t+1} than x_t , leading to rapid growth in consumption as the higher level of resources in $t + 1$ allows for a relaxation of the precautionary saving motive. On the other hand, if the consumer starts with a large value of x_t (greater than x^*), the precautionary motive will be weak and will be outweighed by impatience. The consumer will spend

¹⁵See Carroll (2001a) for derivations of these equations.

more than his expected income, leading (in expectation) to a lower value of x_{t+1} next period, and a lower value of c_{t+1} than c_t ; hence expected consumption growth will be low for large values of x_t .

One might suppose that the level of x_t where the expected growth rate of consumption equals the underlying growth rate of permanent income would be at the target cash-on-hand, x^* . In fact, the figure shows that at x^* , expected consumption growth is slightly lower (by an amount γ) than the growth rate of permanent income. The reason has to do with the concavity of the consumption function, but is not of much intrinsic interest. For purposes of manipulating the diagram, we will just assume γ is a constant, which numerical exercises show is a reasonable approximation.

Assuming γ is constant makes it easy to examine the effects of changing the model's parameters. For example, consider increasing the growth rate to $g' = g + \gamma$ (shown as the dashed horizontal line). If growth is $g + \gamma$, then point at which the $E_t[\Delta \log C_{t+1}]$ curve intersects the *original* g curve will be exactly γ below the g' curve, and thus this intersection will indicate the new target value of cash-on-hand, $x^{**} < x^*$.¹⁶ The new target is at a lower level of cash-on-hand, and (consequently) a higher expected variance of consumption growth. This is simply the human wealth effect in this model: Consumers who expect to have higher income in the future are less willing to save today, so they end up holding a lower buffer stock and suffering a greater degree of consumption variance.

Note a key implication of the figure: Far from being unpredictable *a la* Hall (1978), consumption growth between t and $t+1$ should be related to anything that is related to period- t wealth or income or to the expected variance of consumption growth between t and $t+1$ (such as, for example, the variance of income shocks).

Now consider the implications of this analysis for attempts to detect liquidity constraints by looking for violations of the first-order approximation to the Euler equation, (11). A pioneering paper by Zeldes (1989) pointed out that liquidity constrained consumers would be expected to have faster consumption growth than unconstrained consumers, *ceteris paribus*, because constraints were keeping their consumption lower than they would like. Zeldes's methods for identifying liquidity constrained consumers involved finding households with low levels of assets or current income (the two components of 'cash-on-hand' in the model above) and examining whether such households had faster subsequent consumption growth than others with more cash-on-hand. He found evidence that they did, and concluded that these consumers were liquidity constrained.

But the thrust of the analysis above was that consumers with low wealth or current income should have higher expected consumption growth *even if they are not liquidity constrained*. This illustrates a general principle: The implications of precautionary

¹⁶One more theoretical subtlety: Assuming growth is $g + \gamma$ would also cause changes in the $E_t[(\Delta \log C_{t+1})^2]$ component of the $E_t[\Delta \log C_{t+1}]$ locus, but ignoring these changes (as is done in the diagram) gives the right qualitative answer. For further discussion of this figure, see Carroll (1997).




Figure 3: Expected Consumption Growth and Cash-On-Hand

saving and liquidity constraints for consumption growth are virtually indistinguishable. The reason is that the precautionary motive reduces consumption in precisely the same circumstances as a liquidity constraint would: when cash-on-hand is low. Precautionary saving is in essence like a self-imposed, ‘smoothed’ liquidity constraint. Muddying the waters even further, Zeldes (1984) showed numerically that liquidity constraints can induce a precautionary saving motive even for consumers who have quadratic utility functions and thus do not have a utility-based reason to engage in precautionary saving (see Carroll and Kimball (2001) for a precise statement of the conditions under which liquidity constraints induce precautionary saving, and proofs).

For some purposes, this virtual observational equivalence may not matter much. For example, macroeconomic modellers often need to know what the marginal propensity to consume will be out of some observed transitory fluctuation in aggregate income. Either liquidity constraints or precautionary saving would suggest that some consumers will have a high MPC, and from the modeller’s standpoint it may not matter whether the high MPC results from liquidity constraints or precautionary saving. All that really matters is what fraction of the population is sufficiently impatient that they have a marginal propensity to consume that is much higher than implied by the perfect foresight version of the model.

3.3 Other Methods of Identifying Liquidity Constraints

Of course, there are other purposes for which it is important to distinguish between liquidity constraints and precautionary behavior, most notably in the analysis of the consumption effects of policies that affect credit supply. Fortunately, the fact that it is difficult to distinguish precautionary saving from liquidity constraints using Euler equations does not mean that the two hypotheses cannot be distinguished using other methods. The most promising route is to look at wealth holdings, rather than consumption growth.

The simplest form of liquidity constraint is one in which all borrowing must be collateralized so that consumers are prohibited from having negative net worth. Appending such a constraint to the problem specified above actually has no effect on behavior, since the possibility of the dreaded zero-income-events means that consumers would not have chosen to borrow anyway. However, one could plausibly argue that in modern industrial societies the social safety net prevents consumption from falling all the way to zero, mitigating the impact of unemployment spells. To capture the existence of such a social safety net, suppose that the worst possible event is now defined as an unemployment spell in which income drops to 50 percent of its usual level, an event that occurs with probability $p = 0.05$ to produce a 5 percent aggregate unemployment rate. What does optimal behavior look like with such a social safety net if consumers are prohibited from borrowing?

For baseline values of other parameters, the converged consumption rule is depicted as the locus labelled ‘No Borrowing’ in Figure 4. Below a certain level of cash-on-hand, it is optimal to spend everything, so that the consumption rule coincides with the 45 degree line. Above this cutoff, the consumption function is again concave; since concavity of the consumption function was responsible for most of the insights discussed above (including the endogeneity of consumption growth with respect to the level of wealth and to preference parameters), those insights carry over to the liquidity constrained model for consumers for whom constraints are not currently binding.

A telltale sign of liquidity constraints is visible in the steady-state wealth distribution function, depicted in Figure 5. Whereas the CDF for wealth was completely smooth in the model with precautionary saving but no binding constraints (Figure 2), with constraints there is a mass of households with exactly zero wealth, corresponding to the small vertical segment at the left edge of the CDF. These are the households who were on the 45 degree line portion of $c(x)$ in the previous period and consumed all their resources. Thus, a potential measure of the proportion of the population for whom liquidity constraints are currently binding is simply the proportion for whom wealth (or liquid wealth) is exactly zero.

Panel B. of Table 1 presents the summary statistics for average behavior in the steady-state for this model. The mean and median amount of buffer-stock wealth are both now around 0.25, or about 2 months’ worth less of income than in the unconstrained case. Precautionary savings are lower because the zero income events have

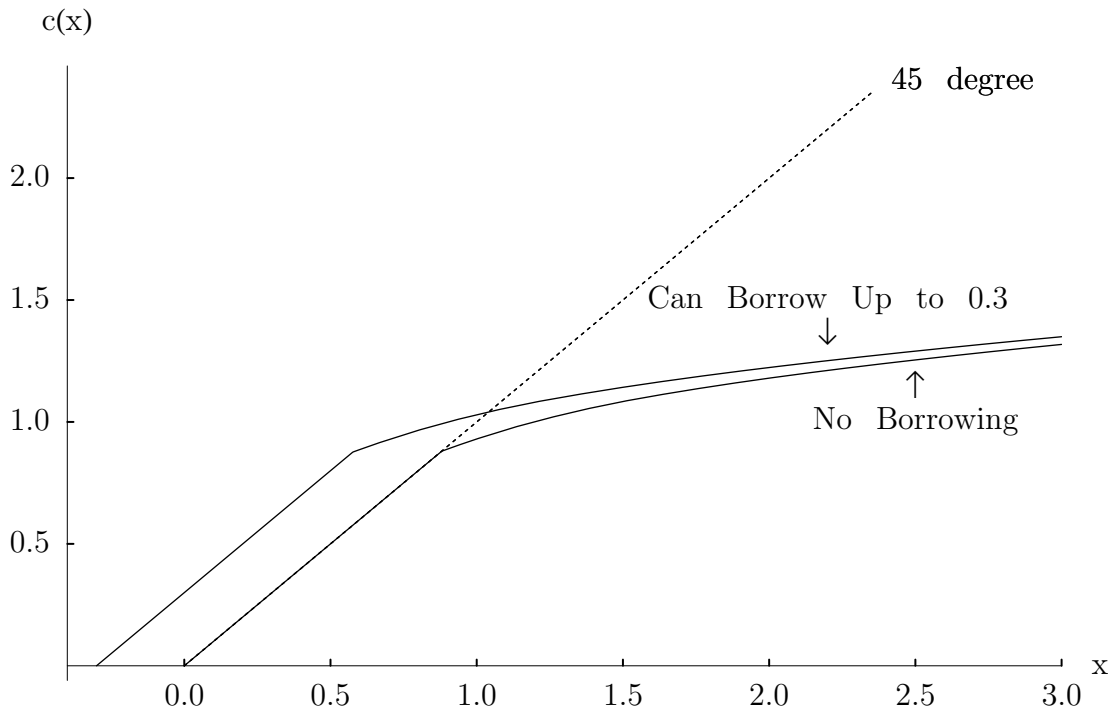


Figure 4: Converged Consumption Rule Under Liquidity Constraints

now been replaced with a comparatively generous unemployment insurance system. Note, however, that the average MPC in the population is roughly the same as under the baseline parameter values in Panel A; furthermore, the effect on the MPC of making consumers more patient is also virtually identical to that in Panel A: for patient consumers, the MPC drops to about 6 percent.

Of course, a complete inability to borrow is unrealistic in modern America, where even household pets receive unsolicited offers of credit cards (and sometimes accept them! see Bennett (1999)). Figures 4 and 5 therefore present the consumption function and steady-state wealth distribution when consumers are allowed to borrow, but only an amount up to thirty percent of their permanent labor income (Ludvigson (1999) presents evidence that actual lenders do strive to limit the ratio of the borrower's debt to income in this manner.) The effect is essentially just to shift the no-borrowing consumption function and CDF to the left by what appears to be about 0.3; Panel C. of Table 1 confirms that mean and median wealth decline by about 0.3. Note that the steady-state average marginal propensity to consume is essentially *the same* as when consumers were prohibited from borrowing. This may go against the grain of intuition, since the natural supposition would seem to be that consumers who can borrow should be better able to shield their consumption against income shocks. But remember that precautionary motives are the only reason these impatient consumers do any saving in the first place. The 'buffering capacity' of a given level of wealth depends on how

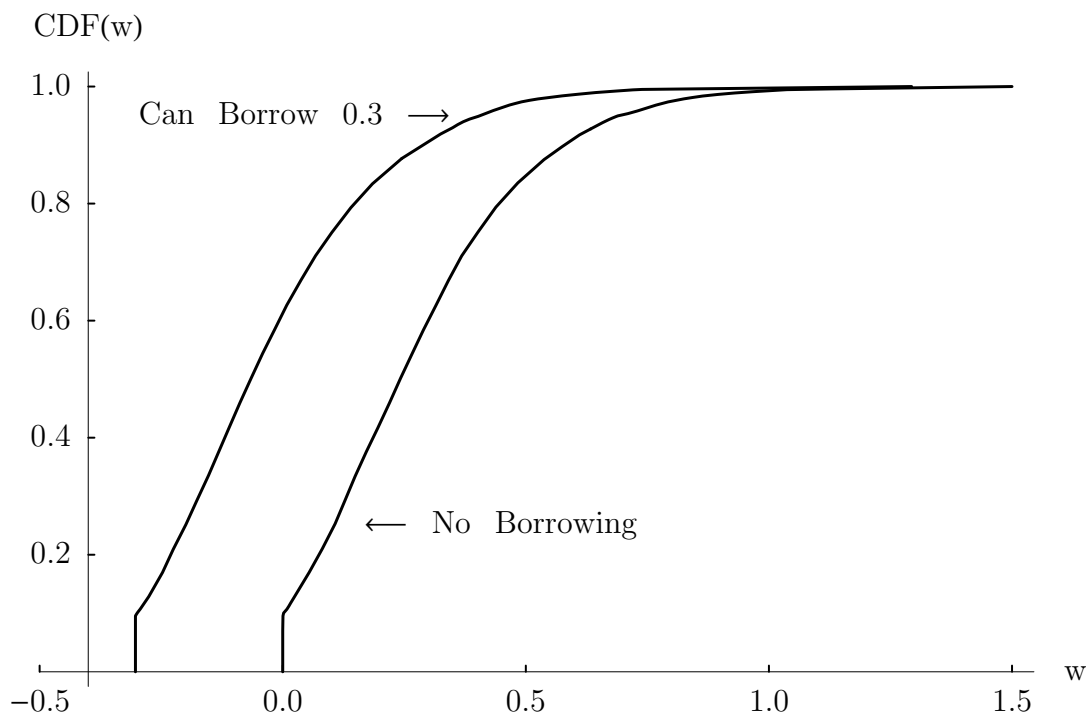


Figure 5: Steady-State Distribution of Wealth with Constraints

much lower wealth could potentially be driven in the case of a bad shock, so allowing borrowing just shifts the whole consumption locus and CDF left, without changing steady-state consumption behavior.

Collectively, the results in Panels A. through C. of the table demonstrate that liquidity constraints are neither necessary nor sufficient to generate a high MPC. What is both necessary and sufficient is impatience, whether there are liquidity constraints or not.

The point that the average MPC depends on impatience rather than the presence or absence of constraints means that many traditional tests of liquidity constraints are questionable at best. For example, Campbell and Mankiw (1991) argue that differences across countries in the sensitivity of consumption growth to predictable income growth may reflect differences in the degree of liquidity constraints, while Jappelli and Pagano (1989) suggest that constraints may be stronger in countries in which consumption growth exhibits excess sensitivity to lagged income growth. It is not clear that either of these interpretations is valid. Instead, the warranted conclusion would seem to be that countries in which consumption exhibits excess sensitivity to lagged or current income may have more households who more impatient, and consequently inhabit the portion of the consumption function where the MPC is high.

If empirical evidence on excess sensitivity of consumption to income is not informative about whether liquidity constraints are important, what kind of evidence would




Figure 6: Empirical CDF of Ratio of Net Worth to Permanent Income, 1995 SCF

be? One example is given by recent work of Gross and Souleles (2000). These authors have managed to obtain a database containing credit report information on a representative sample of consumers, and they show that exogenous increases in households' credit limits result in a substantial increase in actual total debt burdens; in fact, the observed behavior appears to be qualitatively similar to the simulation results presented in Panels B. and C. of the table, in the sense that the debt load after the credit expansion appears to stabilize at a point which provides roughly the same amount of unused credit capacity as before the expansion in the credit line.

Another approach would start with the point, noted above, that the wealth distribution under constraints contains a mass of households at zero wealth (or at the borrowing limit when that is different from zero). For comparison, Figure 6 presents the corresponding cumulative distribution function for data from the 1995 US *Survey of Consumer Finances* on the ratio of nonhousing wealth to permanent income for US consumers between the ages of 25 and 50 (the age range for which the baseline buffer-stock model has been claimed as a plausible description of behavior).¹⁷ Although it is hard to see in the figure, there is indeed a small concentration of households (about 2.5 percent of the population, as indicated in Panel E of Table 1) at exactly the zero-wealth point, and a total of about 10 percent have net worth in the range from zero to two weeks' worth worth (one paycheck) of their permanent income. However, the overall shape of the distribution function (and especially the lower tail) much more closely resembles the shape of the CDF in the unconstrained model, Figure 2, than that in the constrained models, Figure 5; recall also that it is easy to get the unconstrained model to permit negative wealth by assuming a positive minimum value of future income.

The main reason the CDF for the model that allows borrowing fails to match the empirical CDF is that the model implies that there will be a large mass of people who have borrowed up to the maximum credit limit, but the only place in the empirical data where there is any substantial mass is at exactly zero wealth. In fact, Panel C shows that the model predicts essentially zero consumers exactly at zero wealth - because there is nothing special about zero wealth in this model. There is, however, a final element of realism that can be added to the model with constraints that brings its predictions more into accord with the empirical CDF: We can assume that the interest rate at which consumers can borrow is higher than the rate that they can earn on savings. Specifically, if we assume that $R^{\text{borrow}} = 1.15$ (roughly reflecting credit card interest rates in the US), we obtain the consumption function presented in Figure 7. The segment of the new consumption function that lies along the 45 degree

¹⁷Housing and vehicle wealth has been excluded on the grounds that the model does not pretend to be able to capture the complexities associated with durable goods investment. See Carroll and Dunn (1997) for simulation results showing that even when durable goods are added to the model, 'buffer stock' saving behavior emerges with respect to liquid asset holdings. 'Permanent income' is taken to be actual income for the subset of households who said that their income in the survey year was 'about normal.'




Figure 7: Consumption Function with Credit Card Borrowing

line corresponds to the range of x for which the interest rate on saving is not large enough to induce positive saving, but the interest rate on borrowing is high enough to make consumers not want to borrow. At a sufficiently low level of cash-on-hand, however, it becomes worthwhile to borrow even at a 15 percent interest rate, and so the consumption function rises above the 45 degree line.

The CDF from this model is presented in Figure 8. Not only does the model match the bottom tail of the distribution, it also delivers the implication that a small mass of consumers will have exactly zero wealth, just as found in the empirical data.

We now have two models that can match both the high empirical MPC and the general shape of the lower to middle part of the empirical wealth/permanent income ratio. The version without constraints (and with a positive minimum income so that buffer-stock wealth is actual wealth in excess of the lowest possible wealth $-\underline{H}$) has the attraction of simplicity, while the version with liquidity constraints (in the form of both an absolute limit on the amount of borrowing and of differing interest rates for borrowing and lending) has the attraction of greater realism but the cost that it is substantially more complicated, and thus harder to solve.

However, one problem for both models is evident from a closer look at the upper part of the empirical CDF (Figure 6). Although the empirical median wealth/income ratio, at about 0.3, is in the vicinity of the small values predicted by all the models under baseline parameter values, the upper part of the empirical distribution contains vastly




Figure 8: Steady-State Wealth Distribution with Credit Card Borrowing

more wealth than is implied by the model; Panel E. of Table 1 indicates that the mean value of the empirical w is much greater than its median, indicating the skewness of the distribution. Thus, while the presence of substantial numbers of impatient consumers may be essential for reproducing the empirical finding of a high average marginal propensity to consume, the presence of some patient consumers is also required if the model is to match the overall amount of wealth in the US. Whether a life cycle version of the model can match the entire distribution of wealth is a matter of ongoing debate; my own view is that the model certainly cannot match the behavior of the richest few percent in the distribution (unless a bequest motive is added), but may be able to match much of the rest.¹⁸

4 Limitations

I have argued here that the modern version of the dynamically optimizing consumption model is able to match many of the important features of the empirical data on

¹⁸See Huggett (1996), Dynan Skinner and Zeldes (1996), Quadrini and Ríos-Rull (1997), Engen, Gale, and Uccello (1999), and Carroll (2000c) for several perspectives on this question. For general equilibrium macro models which attempt to match both micro and macro data using mixed population of patient and impatient consumers, see Krusell and Smith (1998) and Carroll (2000b).

consumption and saving behavior. There are, however, several remaining reasons for discomfort with the model.

One problem is the spectacular contrast between the sophisticated mathematical apparatus required to solve the optimal consumption problem and the mathematical imbecility of most actual consumers. We can turn, again, to Milton Friedman for a potentially plausible justification for such mathematical modelling. Friedman (1953) argued that repeated experience in attempting to solve difficult problems could build good intuition about the right solution. His example was an experienced pool player who does not know Newtonian mechanics, but has an excellent intuitive grasp of where the balls will go when he hits them. This parable may sound convincing, but some recent work I have done with Todd Allen (2001) suggests that it may sound more convincing than it should. We examine how much experience it would take for a consumer who does not know how to solve dynamic optimization problems to learn nearly optimal consumption behavior by trial and error. Under our baseline setup, we find that it takes about a million ‘years’ of model time to find a reasonably good consumption rule by trial and error. This result may sound preposterous, but we are fairly confident that our qualitative conclusion will hold up, because if there were some trial-and-error method of finding optimal consumption behavior without a large number of trials (and errors), such a method would also constitute a fundamental breakthrough in numerical solution methods for dynamic programming problems. We suspect that the total absence of trial-and-error methods from the literature on optimal solution methods for dynamic optimization problems indicates that such methods are very inefficient, even compared to the enormous computational demands of traditional dynamic programming solution methods. We conclude by speculating that there may be more hope of consumers finding reasonably good rules in a “social learning” context in which one can benefit from the experience of others. However, even the social learning model will probably take considerable time to converge on optimal behavior, so this model provides no reason to suppose that consumers will react optimally in the short- or medium-run to the introduction of new elements into their environment.

As an example of such a change in the consumption and savings environment, consider the introduction of credit cards. In a trial-and-error economy, many consumers would need to try out credit cards, discover that their heavy use can yield lower utility if they lead to high interest payments, and communicate this information to others before there would be any reason to expect the social use of credit cards to approximate their optimizing use. This social learning process could take some time, and even the passage of a recession or two.

There certainly seems to be strong evidence that many American households are now using credit cards in nonoptimal ways. The optimal use of credit cards (at least as implied by solving the final optimizing model discussed above) is as an emergency reserve to be drawn on only rarely, in response to a particularly bad shock or series of shocks. However, the median household with at least one credit card holds about

\$7,000 in debt on all cards combined; that \$7,000 is the balance on which interest is paid, not just the transactions use (Gross and Souleles (2000)). Laibson, Repetto, and Tobacman (1999) argue that this pattern results from time-inconsistent preferences in which consumers have a powerful preference for immediate consumption. Their approach is discussed further in the paper in this symposium.

Another set of empirical findings that are very difficult to reconcile with the modern model of consumption presented here comes in the relationship between saving and income growth, either across countries or across households. A substantial empirical literature has found that much and perhaps most of the strong positive correlation between saving and growth across countries reflects causality from growth to saving rather than the other way around (see Carroll, Overland, and Weil (2000) for a summary). This is problematic because the model implies that consumers expecting faster growth should save less, not more (cf. the model simulations in Table 1). Carroll, Overland, and Weil (2000) suggest that the puzzle can be explained by allowing for habit formation in consumption preferences, but as yet, there is no consensus answer to this puzzle.

A final problem for the standard model is its inability to explain household portfolio choices. The “equity premium puzzle” over which so much ink has been spilled (for a summary see Siegel and Thaler (1997)) remains a puzzle at the microeconomic level, where standard models like the ones presented here imply that consumers should hold almost 100 percent of their wealth in the stock market (for simulation results, see, e.g., Fratantoni (1998), Cocco, Gomes, and Maenhout (1998), Gakidis (1998), Hochguertel (1998), Bertaut and Haliassos (1997)).

5 Conclusion

We shall never cease from exploration
And the end of all our exploring
Will be to arrive where we started
And know the place for the first time.

- T.S. Eliot, “Four Quartets”

Few consumption researchers today would defend the perfect foresight or certainty equivalent models as adequate representations either of the theoretical problem facing consumers or of the actual behavior consumers engage in. Most would probably agree that Milton Friedman’s original intuitive description of behavior was much closer to the mark, at least for the median consumer. It is tempting therefore to dismiss most of the work between Friedman (1957, 1963) and the new computational models of the 1980s and ‘90s as a useless diversion. But a more appropriate view would be that solving and testing those first formal models was an important step on the way to obtaining our current deeper understanding of consumption theory, just as (in a much grander

way) the development of Newtonian physics was a necessary and important precursor to Einstein's general theory.

Understanding of the quantitative implications of the new computational model of consumption behavior is by no means complete. As techniques for solving and simulating models of this kind disseminate, the coming decade promises to produce a flood of interesting work that should define clearly the conditions under which observed consumption, portfolio choice, and other behavior can or cannot be captured by the computational rational optimizing model. Indeed, one purpose of this paper is to encourage readers to join in this enterprise - a process that I hope will be made considerably easier by the availability on the author's website (see the address on the first page) of a set of *Mathematica* programs capable of solving and simulating quite general versions of the computational optimal consumption/saving problem described in this paper.

References

- ALLEN, TODD M., AND CHRISTOPHER D. CARROLL (2001): “Individual Learning About Consumption,” *Forthcoming, Macroeconomic Dynamics*, <http://www.econ.jhu.edu/people/ccarroll/IndivLearningAboutC.pdf> .
- BENNETT, LENNIE (1999): “Platinum Pooch,” *St. Petersburg Times*.
- BERTAUT, CAROL C., AND MICHAEL HALIASSOS (1997): “Precautionary portfolio behavior from a life-cycle perspective,” *Journal of Economic Dynamics And Control*, 21(8-9), 1511–1542.
- BODKIN, RONALD (1959): “Windfall Income and Consumption,” *American Economic Review*, 49(4), 602–614.
- CAGETTI, MARCO (1999): “Wealth Accumulation Over the Life Cycle and Precautionary Savings,” *Manuscript, University of Chicago*.
- CAMPBELL, JOHN Y., AND ANGUS S. DEATON (1989): “Why Is Consumption So Smooth?,” *Review of Economic Studies*, 56, 357–74.
- CAMPBELL, JOHN Y., AND N. GREGORY MANKIW (1991): “The Response of Consumption to Income: A Cross-Country Investigation,” *European Economic Review*, 35, 723–67.
- CARROLL, CHRISTOPHER D. (1992): “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence,” *Brookings Papers on Economic Activity*, 1992(2), 61–156.
- (1994): “How Does Future Income Affect Current Consumption?,” *The Quarterly Journal of Economics*, CIX(1), 111–148, <http://www.econ.jhu.edu/people/ccarroll/howdoesfuture.pdf> .
- (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, CXII(1), 1–56, <http://www.econ.jhu.edu/people/ccarroll/BSLCPIH.pdf> .
- (2000a): “Lecture Notes on Solving Microeconomic Dynamic Stochastic Optimization Problems,” <http://www.econ.jhu.edu/people/ccarroll>.
- (2000b): “Requiem for the Representative Consumer? Aggregate Implications of Microeconomic Consumption Behavior,” *American Economic Review, Papers and Proceedings*, pp. 110–115, <http://www.econ.jhu.edu/people/ccarroll/RequiemFull.pdf>, Table 1 Data: <http://www.econ.jhu.edu/people/ccarroll/requiemfiles.zip> .

- (2000c): “Why Do the Rich Save So Much?,” in *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, ed. by Joel B. Slemrod, pp. 463–485. Harvard University Press, <http://www.econ.jhu.edu/people/ccarroll/Why.pdf> .
- (2001a): “Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second Order Approximation),” *Advances in Macroeconomics*, 1(1), <http://www.bepress.com/bejm/advances/vol1/iss1/art6> .
- (2001b): “Theoretical Foundations of Buffer Stock Saving,” *Manuscript, Johns Hopkins University, June*, <http://www.econ.jhu.edu/people/ccarroll/BufferStockTheory.pdf> .
- CARROLL, CHRISTOPHER D., AND WENDY E. DUNN (1997): “Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets,” in *NBER Macroeconomics Annual, 1997*, ed. by Benjamin S. Bernanke, and Julio Rotemberg, pp. 165–229. MIT Press, Cambridge, MA, <http://www.econ.jhu.edu/people/ccarroll/macroann.pdf>; Methodological Appendix: <http://www.econ.jhu.edu/people/ccarroll/methods3.pdf>; Empirical Results and Simulation Programs: <http://www.econ.jhu.edu/people/ccarroll/cdfiles.html>; .
- CARROLL, CHRISTOPHER D., AND MILES S. KIMBALL (1996): “On the Concavity of the Consumption Function,” *Econometrica*, 64(4), 981–992, <http://www.econ.jhu.edu/people/ccarroll/concavity.pdf> .
- (2001): “Liquidity Constraints and Precautionary Saving,” *Manuscript, Johns Hopkins University*.
- CARROLL, CHRISTOPHER D., JODY R. OVERLAND, AND DAVID N. WEIL (2000): “Saving and Growth with Habit Formation,” *American Economic Review*, 90(3), 341–355, <http://www.econ.jhu.edu/people/ccarroll/AERHabits.pdf> .
- COCCO, JOAO, FRANCISCO J. GOMES, AND PASCAL J. MAENHOUT (1998): “Consumption and Portfolio Choice Over the Life Cycle,” *Manuscript, Harvard University*.
- DEATON, ANGUS S. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59, 1221–1248.
- DYNAN, KAREN E., JONATHAN S. SKINNER, AND STEPHEN P. ZELDES (1996): “Do the Rich Save More?,” *Manuscript, Board of Governors of the Federal Reserve System*.
- ENGEN, ERIC, WILLIAM GALE, AND CORI UCCELLO (1999): “The Adequacy of Retirement Saving,” *Brookings Papers on Economic Activity*, 1999(2).

- FRATANTONI, MICHAEL C. (1998): "Income Uncertainty and the Equity Premium Puzzle," *Manuscript, Johns Hopkins University*.
- FRIEDMAN, MILTON A. (1953): *Essays in Positive Economics*. University of Chicago Press, Chicago.
- (1957): *A Theory of the Consumption Function*. Princeton University Press.
- (1963): "Windfalls, the 'Horizon,' and Related Concepts in the Permanent Income Hypothesis," in *Measurement in Economics*, ed. by Carl Christ, et al., pp. 1–28. Stanford University Press.
- GAKIDIS, HARALOBOS E. (1998): "Stocks for the Old? Earnings Uncertainty and Life-Cycle Portfolio Choice," *Manuscript, Massachusetts Institute of Technology*.
- GOURINCHAS, PIERRE-OLIVIER, AND JONATHAN PARKER (1999): "Consumption Over the Life Cycle," *Manuscript, Princeton University*.
- GROSS, DAVID B., AND NICHOLAS S. SOULELES (2000): "Consumer Response to Credit Supp[ly]: Evidence from Credit Card Data," *Manuscript, University of Pennsylvania*.
- HALL, ROBERT E. (1978): "Stochastic Implications of the Life-Cycle/Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 96, 971–87.
- HALL, ROBERT E., AND FREDERIC MISHKIN (1982): "The Sensitivity of Consumption to Transitory Income: Evidence from PSID Households," *Econometrica*, L, 461–81.
- HOCHGUERTEL, STEFAN (1998): "A Buffer Stock Model with Portfolio Choice: Implications of Income Risk and Liquidity Constraints," *Manuscript, Uppsala University*.
- HUGGETT, MARK (1996): "Wealth Distribution in Life Cycle Economies," *Journal of Monetary Economics*, 38(3), 469–494.
- KEYNES, JOHN MAYNARD (1935): *The General Theory of Employment, Interest, and Money*. Harvest/HBJ, San Diego, New York, London, 1964.
- KREININ, MORDECAI E. (1961): "Windfall Income and Consumption: Additional Evidence," *American Economic Review*, 51, 388–390.
- KRUSELL, PER, AND ANTHONY A. SMITH (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106(5), 867–896.
- LAIBSON, DAVID, ANDREA REPETTO, AND JEREMY TOBACMAN (1999): "A Debt Puzzle," *Manuscript, Harvard University*.

- LANDSBERGER, M. (1966): “Windfall Income and Consumption: Comment,” *American Economic Review*, 56, 534–540.
- LUDVIGSON, SYDNEY (1999): “Consumption and Credit: A Model of Time-Varying Liquidity Constraints,” *The Review of Economics and Statistics*, 81(3).
- MAYER, THOMAS (1972): *Permanent Income, Wealth, and Consumption*. University of California Press, Berkeley.
- MERTON, ROBERT C. (1969): “Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case,” *Review of Economics and Statistics*, 50, 247–257.
- MODIGLIANI, FRANCO (1966): “The Life Cycle Hypothesis of Saving, the Demand for Wealth, and the Supply of Capital,” *Social Research*, 33, 160–217.
- MODIGLIANI, FRANCO, AND RICHARD BRUMBERG (1954): “Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data,” in *Post-Keynesian Economics*, ed. by Kenneth K. Kurihara, pp. 388–436. Rutgers University Press, New Brunswick, N.J.
- QUADRINI, VINCENZO, AND JOSÉ-VÍCTOR RÍOS-RULL (1997): “Models of the Distribution of Wealth,” *Manuscript, University of Pennsylvania*.
- SAMUELSON, PAUL A. (1969): “Lifetime Portfolio Selection By Dynamic Stochastic Programming,” *Review of Economics and Statistics*, 51, 239–46.
- SAMWICK, ANDREW A. (1995): “The Limited Offset Between Pension Wealth and Other Private Wealth: Implications of Buffer-Stock Saving,” *Manuscript, Department of Economics, Dartmouth College*.
- SIEGEL, JEREMY J., AND RICHARD H. THALER (1997): “Anomalies: The Equity Premium Puzzle,” *Journal of Economic Perspectives*, 11(1), 191–200.
- TULLIO JAPPELLI, MARCO PAGANO (1989): “Consumption and Capital Market Imperfections: An International Comparison,” *The American Economic Review*, 79(5), 1088–1105.
- VIARD, ALAN (1993): “The Productivity Slowdown and the Savings Shortfall: A Challenge to the Permanent Income Hypothesis,” *Economic Inquiry*, 31, 549–564.
- ZELDES, STEPHEN P. (1984): “Optimal Consumption with Stochastic Income,” Ph.D. thesis, Massachusetts Institute of Technology.
- (1989): “Consumption and Liquidity Constraints: An Empirical Investigation,” *Journal of Political Economy*, 97, 305–46.