



Working Paper Series

Consumption-Based Modeling of Long-Horizon Returns

Kent D. Daniel and David A. Marshall

Working Papers Series
Research Department
(WP-98-18)

Federal Reserve Bank of Chicago

December 24, 1998

Consumption-Based Modeling of Long-Horizon Returns[†]

Kent D. Daniel* and David A. Marshall**

Abstract

Numerous studies have documented the failure of consumption-based pricing models to explain observed patterns in stock and bond returns. This failure has sometimes been attributed to frictions, transaction costs or durability. If such frictions are important, they should primarily affect the higher frequency components of asset returns. The long-swings, or lower-frequency comovements should be less affected. Consequently if transaction costs are important, tests of the consumption based asset pricing model which concentrate on lower-frequency components may be more successful.

We investigate this hypothesis using a variety of diagnostic tests. We first use coherence analysis and bandpass filtering analysis to show that, while there is a complete lack of correlation between asset returns and consumption growth at frequencies higher than about 0.7 years^{-1} (swings longer than 1.4 years), the coherence/correlation between the two series at lower frequencies is above 60%. We perform Hansen and Jagannathan (1991) bounds tests, χ^2 tests of moment restrictions, and Hansen and Jagannathan (1997) specification tests of three consumption-based models of the asset-pricing kernel: Time-separable preferences with power utility, the Abel (1990) “Catching up with the Joneses” preferences, and Constantinides (1990) habit-formation preferences. While none of these models perform well at the quarterly horizon, the performance of the Abel and Constantinides models improves strikingly at the two-year horizon.

[†]We thank Steve Cecchetti, Larry Christiano, John Cochrane, George Constantinides, Larry Epstein, Lars Hansen Walter Torous and participants of the NBER Asset Pricing Group Meeting and UCLA Finance Seminar for helpful discussions, and we acknowledge Denise Duffy and Glenn McAfee for superlative research assistance. Daniel gratefully acknowledges research support from the Center for Research in Security Prices (CRSP) at the University of Chicago. The opinions expressed in this paper are those of the authors, and do not reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

*Kellogg Graduate School of Management, Northwestern University, Evanston, IL 60208, 847-491-4113, kent@kent.kellogg.nwu.edu, <http://kent.kellogg.nwu.edu/>

**Federal Reserve Bank of Chicago, 312-322-5102, dmarshall@frbchi.org

1 Introduction

Most research on consumption-based asset pricing focuses on short-horizon returns. The starting point is the familiar intertemporal Euler equation

$$1 = E_t [m_{t+\tau}^\tau r_{t+\tau}^\tau] \quad (1)$$

where $r_{t+\tau}^\tau$ denotes the gross real cumulative asset return from date t to date $t + \tau$ and $m_{t+\tau}^\tau$ denotes the intertemporal marginal rate of substitution (IMRS) between wealth at date t and wealth at date $t + \tau$. The unconditional version of this Euler equation is

$$1 = E [m_{t+\tau}^\tau r_{t+\tau}^\tau] . \quad (2)$$

In the literature on consumption-based asset pricing, m_t^τ is modeled as a function of aggregate consumption, and the implications of (1) or (2) are tested for return horizons τ equal to one month or one quarter. The literature typically rejects the conditional Euler equation (1).

Presumably, the unconditional Euler equation (2) should be easier to fit to the data. Equation (2) only requires that the IMRS correctly price equity returns "on average". Yet, Cochrane and Hansen (1992) show how difficult it is to even fit this unconditional Euler equation. The problem is that, to fit equation (2), one needs substantial covariation between $m_{t+\tau}^\tau$ and $r_{t+\tau}^\tau$. Formally, equation (2) implies

$$E [r_{t+\tau}^\tau - r f_t^\tau] E \left[\frac{1}{r f_t^\tau} \right] = -cov (m_{t+\tau}^\tau, r_{t+\tau}^\tau) = -\sigma_m \sigma_r \rho_{m,r} \quad (3)$$

where $r f_t^\tau$ denotes the gross real risk-free rate from date t to date $t + \tau$, σ_m and σ_r denote the standard deviations of $m_{t+\tau}^\tau$ and $r_{t+\tau}^\tau$ respectively, the correlation coefficient between $m_{t+\tau}^\tau$ and $r_{t+\tau}^\tau$ is denoted $\rho_{m,r}$, and we use the implication of equation (1) that $E[m_{t+\tau}^\tau] = E[1/r f_t^\tau]$. The left-hand side of equation (3) is the mean equity premium, "discounted" by the mean inverse risk-free rate. To generate a large mean equity premium, a model of $m_{t+\tau}^\tau$ must display substantial negative covariation with equity returns. Within the context of consumption-based asset pricing, this would require substantial positive covariation between equity returns and consumption growth.

Empirically, the contemporaneous correlation between quarterly returns and quarterly consumption growth is small (less than 0.15), and the largest correlation at any lead/lag (when

returns lead consumption growth by two quarters) is less than 0.2.¹ Cochrane and Hansen (1992) call this low correlation between the return on market proxies and consumption growth the “correlation puzzle.”² A number of factors have been proposed to account for the low correlations between stock returns and aggregate consumption growth at short-horizons, including uninsurable cross-sectional heterogeneity,³ fixed costs of adjusting consumption,⁴ costs of portfolio adjustment,⁵ and even small deviations from perfect rationality.⁶ While these factors could substantially affect the co-movements of asset returns and aggregate consumption at high frequencies, they should be less disruptive to the theory at longer horizons. Simple correlations between consumption growth and the VW index return suggest that there may be merit in this argument. While the contemporaneous correlation between consumption growth and returns at the one-year horizon is about the same as for the one-quarter horizon, the correlation between one-year consumption growth and one-year returns lagged by two quarters is 0.35, almost twice as high as the maximal correlation found between quarterly returns and quarterly consumption growth

If consumption-based pricing operators perform better at pricing long-horizon assets, this would provide indirect evidence that the basic intuition underlying the equilibrium asset pricing theories of Lucas (1978), Breeden (1979), and Grossman and Shiller (1982) is sound; the well-documented failures of this theory at matching high-frequency data could be attributed to transaction costs, market imperfections, and uninsurable heterogeneity. This analysis would not be a substitute for formal modeling of these frictions, but would at least suggest that this approach is on the right track. If, however, consumption-based pricing proves useless at all frequencies, this would represent a substantial challenge to equilibrium pricing theory. In all such models, financial assets are vehicles for transferring consumption across time and

¹These numbers were calculated using real non-durable and services consumption data, and returns on the CRSP value-weighted index, deflated by the ND&S deflator.

²Even if the correlation between consumption growth and equity returns is small, equation (3) suggests an alternative way of generating a large mean equity premium: increase the variability of the IMRS (for example, by assuming a high degree of risk aversion). Cecchetti, Lam, and Mark (1993) try this second strategy. While they can match the mean equity premium, they have difficulty matching both the first and second moments of the equity and risk-free return data.

³See, e.g., Constantinides and Duffie (1996), Lucas (1994), Heaton and Lucas (1995)

⁴Grossman and Laroque (1990), Marshall (1994), Marshall and Parekh (1998).

⁵Luttmer (1996), He and Modest (1995).

⁶See Cochrane (1989).

random states. It is not clear what sort of model can capture this intuition without implying a theoretical linkage between aggregate consumption and asset returns at some horizon.

In this study, we investigate the ability of consumption growth to explain asset returns at low frequencies using a variety of diagnostic tests. We first use multivariate spectral analysis to characterize the co-movements of consumption growth and excess equity returns. We find that, while the coherence between these variables is insignificant at high frequencies (above 0.7 years⁻¹), at lower frequencies this coherence is high and statistically significant. We confirm the coherence analysis using bandpass filtering analysis similar to that suggested by Baxter and King (1994). We find no statistically significant correlation between the two series at high frequencies. However, at lower “business-cycle” frequencies, we find a correlation of over 50%.

We then turn to more formal tests of consumption-based asset pricing at longer horizons. We look at standard time-separable power utility, the Abel (1990) “Catching up with the Joneses” preferences, and the Constantinides (1990) habit-formation preferences. Our investigation uses three diagnostic tools: first, we apply the Hansen and Jagannathan (1991) mean-variance analysis, modified to take into account the unconditional correlation between the pricing kernel and asset returns; second, we apply the standard χ^2 of the moment restriction in (2). The third diagnostic we use follows Hansen and Jagannathan (1997). They note that some preference specifications may result in an extremely volatile m_t , but one which does not really “fit” the data better. However, because of the increased volatility, such a preference specification may still yield a low χ^2 . Hansen and Jagannathan (1997) and Hansen, Heaton, and Luttmer (1995) suggest a specification test which is immune to this problem. We employ this test as a further diagnostic tool.

According to our empirical results, none of the consumption-based models performs well at the quarterly horizon. However, when the horizon is lengthened to two years, versions of the Abel (1990) and Constantinides (1990) models of preferences perform fairly well. In particular, no significant violations of the Hansen-Jagannathan restrictions are found, the χ^2 statistics are insignificantly different from zero, and the Hansen and Jagannathan (1997) diagnostic reveals rather small discrepancies from equation (2).

Other papers in the literature have asked whether the restrictions of economic theory are

better satisfied with longer-horizon returns than with monthly or quarterly returns. Brainard, Shapiro, and Shoven (1991) find that stock returns correspond more closely to the theoretical return to invested capital at longer horizons. Cochrane and Hansen (1992), however, find that the performance of time-separable preferences deteriorates as the horizon lengthens. Finally, Daniel and Marshall (1997) use a vector ARCH model to estimate the conditional covariance between equity returns and the pricing kernel implied by various models. All models do poorly at the quarterly horizon. At the two-year horizon, however, they find that the mean and variance of this conditional covariance series approximate the corresponding moments of the conditional equity premium as long as sufficient time-nonseparability is incorporated into preferences. Furthermore, the variation through time in this conditional equity premium is to some extent reflected in variation in the conditional covariance series.

The conditional evidence is important. Although we find evidence of a strong unconditional relationship here, the high unconditional correlations between consumption growth and asset returns at business cycle frequencies could be caused by co-movement of the expected growth and return series as opposed to co-movements of the innovations. There is strong evidence of variation in expected returns across the business cycle.⁷ Thus a test of the conditional model might fail when the business cycle variables are used as instruments, even though the unconditional test does not reject the model. The results in Daniel and Marshall (1997) suggest that this should not be the case.

The remainder of the paper is organized as follows: In section 2, we present evidence that although consumption-growth is virtually uncorrelated with asset returns at high frequencies, it displays substantial covariation with asset returns at business-cycle frequencies. In section 3 we characterize the preference specifications to be explored in the remainder of the paper; we discuss the model diagnostics to be used; and we evaluate particular parameterizations of these models at various horizons. In that section, we also discuss problems in modeling long horizon returns that were pointed out by Cochrane and Hansen (1992). We argue that these problems can be resolved, in principle, by the use of time-nonseparable preferences as advocated by Abel

⁷For example, Fama and French (1989) show that the term spread and the default spread forecast future long horizon returns, and Chen (1991) shows that these same variables forecast future economic activity. If these same variables forecast future consumption growth, this could result in an unconditional correlation, even if the conditional correlation is zero.

(1990) and Constantinides (1990). In section 4, we provide a more formal analysis of the Abel and Constantinides models in which the model parameters are chosen to minimize the Hansen and Jagannathan (1997) specification criterion. Section 5 summarizes.

2 Low-Frequency Correlation of Asset Returns and Consumption Growth

In this section, we ask whether the Cochrane and Hansen (1992) correlation puzzle is a phenomenon at *all* frequencies, or just at high frequencies. We use two techniques to explore this question: multivariate spectral or coherence analysis, and an analysis of the correlations of bandpass filtered consumption growth and asset returns.


2.1 Coherence Analysis

To assess the frequency-by-frequency breakdown of the correlation between consumption growth and a market proxy, we perform coherence analysis of the consumption growth and excess VW index return series. Essentially, what the coherence analysis does is to split each of the two series up into a set of Fourier components at different frequencies, and then to determine the correlation of a set of Fourier components for the two series around each frequency. The method we use for generating the coherence series is described in the appendix. In addition, the method yields the phase relation between the two series, which is a measure of how far the series must be shifted to maximize the correlation of the sets of Fourier components.

In Figure 1, the upper panel gives the coherence between the two series, as a function of the frequency. Approximate 95% confidence intervals for both coherence and phase are plotted as dashed lines in the figure.⁸ We also plot the 95% confidence interval that the coherence is above zero, which is the dot-dashed line in the figure.

Figure 1 shows that the coherence is low at high frequencies. However, for lower frequencies, but not the very lowest, the coherences are relatively large, and are considerably more than two standard errors away from zero. This suggests that any correlation that arises between

⁸These are calculated following Bloomfield (1976). For the coherence series, the standard error is inversely proportional to the level of the coherence, and in calculating these confidence intervals we have assumed that the true coherence is equal to the estimator.

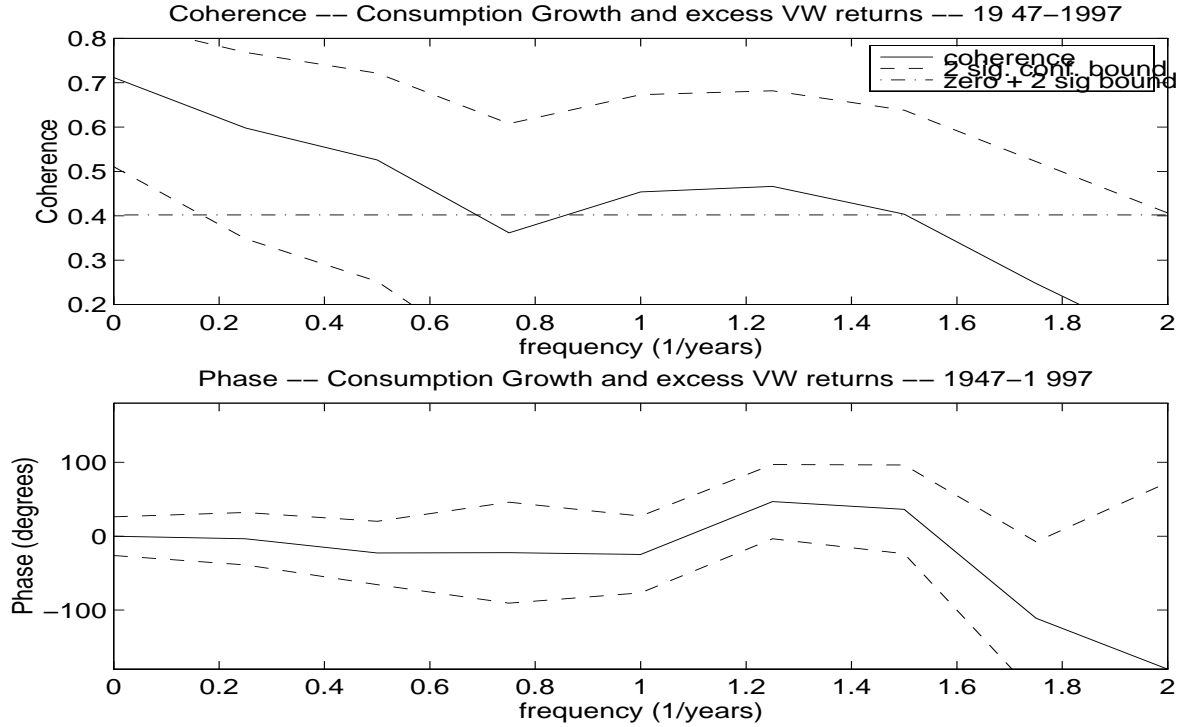


This figure shows, in the upper panel, the coherence between quarterly real non-durable and services consumption growth and quarterly real returns on the CRSP VW index over the period 1947:1-1994:1. The lower panel shows the phase relationship, in degrees, between the two series at the stated frequency. A hamming window and a smoothing parameter (n/m) of 16 is used. Details of the coherence and phase calculations are given in the Appendix.

Figure 1: Coherence and Phase plots for consumption Growth and excess VW index Returns

the consumption growth and asset return series is due to the co-movement at these lower frequencies.

The lower panel of Figure 1 displays the phase as a function of frequency. Notice that at a frequency of zero, the phase is zero, and then decreases approximately linearly with frequency up to a frequency of about 1 year^{-1} . This suggests a constant-length lead/lag relationship between the two variables, since the approximate lead/lag length equals the phase multiplied by $(1/\text{frequency} * 360)$. Performing this calculation, Figure 1 tells us that return series leads the consumption growth series by approximately two quarters. To verify this, in Figure 2 we lag the VW return series by 2 quarters and rerun the coherence analysis. The coherence values increase slightly at most frequencies, except at frequency zero where it increases dramatically. The phase is now approximately zero for frequencies less than one year^{-1} .




The construction of this figure was identical to that for Figure 1, except that the VW index return was lagged by two quarters before calculating the coherence.

Figure 2: **Coherence and Phase plots for Consumption Growth and Lagged excess VW index Returns**

In summary, the coherence analysis suggests while there is little relation between consumption growth and asset returns at high frequencies, there is a strong relationship at low frequencies, but that the relationship is not contemporaneous: consumption growth lags the market return by about two quarters. This might tie in well with a frictions story: If there is an extra cost of adjusting consumption quickly rather than slowly, then this is exactly the sort of relationship one might expect to observe. Alternatively, if agents' preferences exhibited habit formation (e.g., Constantinides (1990)) and there were a utility cost to rapid consumption adjustment, consumption would respond slowly to change in asset prices.

Figure 3: **Bandpass Filter Frequency Response**

This figure shows the gain of the three filters used in the bandpass filtering analysis over the frequency range from 0 to 1 years⁻¹. The gain from 1 to 2 years⁻¹ is not shown, but is approximately 1 for the high filter and 0 for the other two.



2.2 Bandpass Filtering Analysis

To verify the coherence results, we also conduct a time-series bandpass filtering analysis, in which we break the consumption growth and asset returns series into different components using a set of moving-average (MA) filters. To do this, we utilize a bandpass filtering analysis similar to that suggested by Baxter and King (1994), and utilized by Baxter (1994).

The MA filters we utilize here are each designed to pick out a certain range of frequencies in the data. The filters are symmetric two-sided filters, and consequently do not introduce phase distortions into the data. Though the design of the filters uses spectral analysis techniques, the filters themselves are simple MA filters. The process of filtering a data series is conducted by simply convolving the data series with the set of filter weights, as with any MA filter.

The filter design is done using the FIRLS (finite impulse response least squares) function in *MATLAB*.⁹ This function takes as an input a desired piecewise linear spectral response function, and then finds the length n (here 31 quarters) MA filter which provides a spectral response which is as close as possible to the desired spectral shape (in a weighted least-squares

⁹This function is taken from *MATLAB*'s the Signal Processing Toolbox.

sense).

We constructed three filters. To conform to the business cycle literature, we call these the *trend*, *cycle* and *high* filters. The *trend* filter is designed to eliminate all swings in the data shorter than 8 years (*i.e.*, with a frequency of higher than 0.125 years^{-1}). The *cycle* filter is designed to pass all components with wavelength between 1.5 and 8 years, and the *high* filter is designed to eliminate all swings longer than 1.5 years. The spectral gain obtained for the three filters are shown in Figure 3.¹⁰ While we would like these filters to have a gain of 1 over the bandpass region and zero elsewhere, this is not achievable with a finite length filter.¹¹ Instead, this figure shows that the filter frequency responses are about as good as is achievable, given the constraints.

The motivation for eliminating high frequencies from the data has already been discussed. The motivation for eliminating the correlations at very low frequencies is that there may be long-run changes in the structure of the economy that affect the co-movement of consumption and returns on financial assets. (Changes in the legal and institutional framework of the pension-fund industry is one example.) At a more basic level, since other macroeconomic variables move together at business cycle frequencies, and not elsewhere, this may also be true for consumption.

The correlations between various measures of consumption growth and VW return series, filtered using each of the three filters in turn, are presented in Table 1. We consider six alternative measures of consumption: consumer durables, consumer non-durables, consumer services, nondurables plus services, consumer services excluding the implied service flow from owner-occupied housing, and nondurables plus services excluding the service flow from owner-occupied housing. The procedure for constructing each of these correlations was to filter each of the two series, and then calculate the maximum correlation $\hat{\rho}_{max}$ between the two series over a lead-lag range of eight quarters, that is

¹⁰We also demean the set of MA coefficients output from the FIRLS procedure for the *cycle* and *high* filter weight series. Since these are finite length filters, if the sum of the MA coefficients is zero then the filters will eliminate any integrated ($I(1)$) component present in the data.

¹¹Baxter and King (1994) discuss the reasons that an “ideal spectral shape” is not achievable in practice.

Table 1: The Correlation of Bandpass Filtered Returns and Consumption Growth Series

The quarterly real CRSP VW index returns and the real non-durable and services consumption growth series (1947:1-1997:4), were each filtered using one of the three filters, as described in the text. The maximum correlation coefficient between the two filtered series, $\hat{\rho}_{max}$, is computed as in equation (4). The value of $\hat{\rho}_{max}$ as well as the maximizing value of τ are given in the Table for each filter. The Monte-Carlo determined p-values are the probability of obtaining a number at least as large as the sample value assuming that the returns are *i.i.d.* normal, and independent of consumption growth. Six measures of consumption are used: consumer durables ("Dur."); consumer non-durables ("ND"); consumer services ("Serv."); nondurables plus services ("ND&S"); consumer services excluding the implied service flow from owner-occupied housing ("SxH"); and nondurables plus services excluding the implied service flow from owner-occupied housing ("ND&SxH").

		Filter		
		trend	cycle	high
Dur.	$\hat{\rho}_{max}$	-0.1053 (0.718)	0.5343 (0.002)	0.2053 (0.192)
	τ_{max}	0	2	-3
ND	$\hat{\rho}_{max}$	-0.0521 (0.657)	0.5092 (0.003)	0.2735 (0.019)
	τ_{max}	5	2	1
Serv.	$\hat{\rho}_{max}$	-0.0148 (0.619)	0.4733 (0.020)	0.1449 (0.634)
	τ_{max}	8	2	4
ND&S	$\hat{\rho}_{max}$	-0.0884 (0.703)	0.5497 (0.002)	0.1662 (0.443)
	τ_{max}	8	2	1
SxH	$\hat{\rho}_{max}$	-0.1427 (0.747)	0.4248 (0.052)	0.1344 (0.698)
	τ_{max}	8	2	4
ND&SxH	$\hat{\rho}_{max}$	-0.1504	0.5360 (0.002)	0.1739 (0.38)
	τ_{max}	8	2	1

$$\hat{\rho}_{max} = \max_{\tau} \left(\frac{1}{\sigma(r)\sigma(\Delta c)} \right) \left(\frac{1}{T - \tau} \right) \sum_{t=\max\{1, 1+\tau\}}^{\min\{T, T+\tau\}} (r_{t-\tau} \cdot \Delta c_t) \quad -8 \leq \tau \leq 8. \quad (4)$$

where r_t and Δc_t are demeaned returns and consumption growth over period t, respectively, and $\sigma(r)$ $\sigma(\Delta c)$ are the corresponding sample standard deviations for the series.

We use Monte-Carlo analysis to determine the significance levels of the correlations in Table 1 both in order to account for the fact that we are reporting the maximum of 17 coefficients, and to account for the serial correlation induced in the two series by the filtering operations. The *p-values* reported in Table 1 are therefore the probabilities of obtaining a coefficient at least as large as the one found if the quarterly returns series were *i.i.d.* normal, and independent of consumption growth.

As predicted by the coherence analysis, the maximum correlation between the two series after filtering using the cycle filter is high. For example, the maximized correlation coefficient is 54.97% for ND&S consumption, with a one-tailed p-value of 0.2%. Also, as predicted by the coherence analysis, the correlation is maximized when the returns series is lagged by two quarters.

We find little correlation between the returns series and any of the high-pass filtered consumption growth series. The maximized correlation coefficient is about 20% for most measures of consumption, and is insignificantly different from zero. An exception is the non-durable series, which is significant at the 1.9% (one-tailed) value. To test whether the correlation is different at the high and business cycle frequencies, we conduct an additional Monte-Carlo simulation. Note that the correlation between the business-cycle filtered series is approximately 0.5 for all measures of consumption. In our simulation exercise, we impose the null hypothesis that the correlation between the high-pass filtered series is also 0.5 in population. We then compute the probability of obtaining a maximum observed correlation of 0.27 (the highest correlation in the last column of Table 1). We find that this probability is less than 0.001. We conclude that the correlations between high-pass filtered series are significantly less than the correlations between business-cycle filtered series.

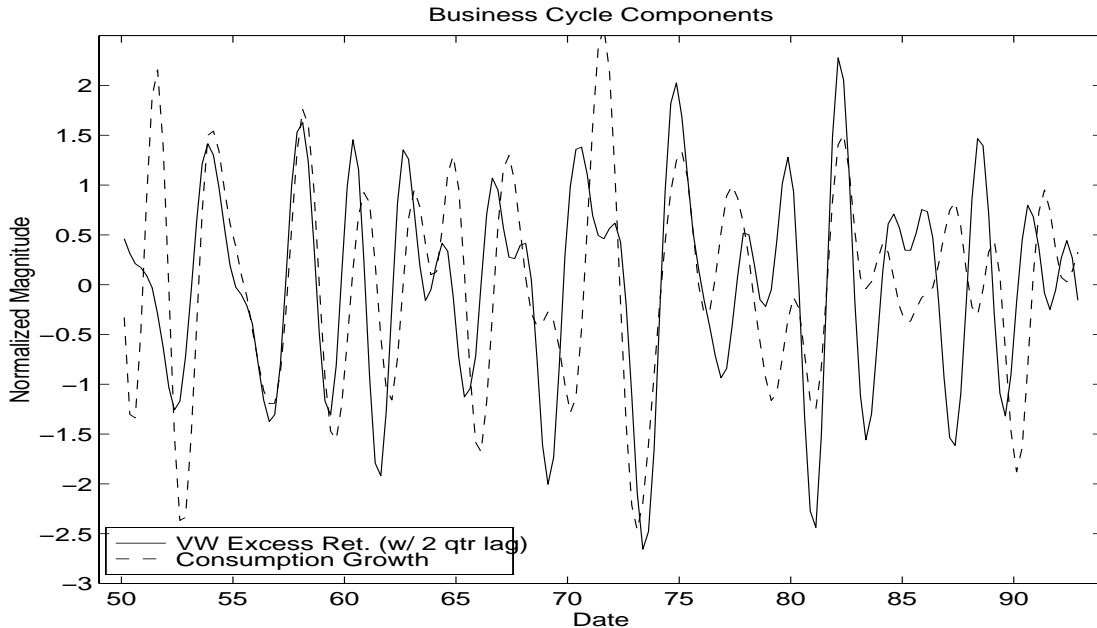
Interestingly, the maximum correlation of the trend components of the two series is also insignificant, at -8.8% (one-tailed p-value of 70.3%) for non-durable and services. A Monte-Carlo simulation shows that the probability of getting a measured trend-correlation of less than zero is 4.8%, when the true correlation between the two series equal to 0.5 at all frequencies. Thus, it appears that the correlation is indeed lower at trend frequencies than at business cycle frequencies.

The normalized, *cycle*-filtered asset returns and consumption growth series are plotted in Figure 4.¹² We have also lagged the VW index return series by two quarters. The plot shows the very strong relationship between the two series in this frequency range from the 50's through the early 80's. The late 80's are the only period in the plot where the relation is not extremely strong, perhaps because of the 1987 crash.

¹²To normalize the series, we divide each series by its sample standard deviation.

Figure 4: **Filtered Asset Returns and Consumption Growth Series**

This figure plots the business-cycle filtered real consumption growth (Total) and filtered, two-quarter-lagged, VW returns series. The filtering method is described in the text. The period is 1947:1-1994:1, but the first and last four years are truncated in the filtering process. Each series is normalized by dividing by its standard deviation.



3 Unconditional Moments Tests of the Asset Pricing Relation

Section 2 documents that consumption growth is more highly correlated with equity returns at business cycle horizons than at short horizons, and indeed that there is insignificant correlation at frequencies of less than 0.67 years^{-1} . This suggests that there may be frictions which “de-link” consumption and asset price movements at high frequencies. If so, tests of the consumption based asset pricing model may be considerably more successful at pricing longer-horizon returns. For the remainder of the paper, we explore this conjecture at various horizons.

3.1 Preference Models

3.1.1 Time-Separable Preferences

The most widely-studied (and widely-rejected) preference specification in the consumption-based pricing literature is time-separable power utility. In this specification, agents solve the

following maximization problem:

$$\max_{\{c_{t+j}\}_{j=0}^{\infty}} U \equiv E_t \sum_{j=0}^{\infty} \beta^j \frac{c_{t+j}^{1-\gamma}}{1-\gamma} \quad (5)$$

subject to the usual budget constraint. The τ -period IMRS is:

$$m_t^\tau = \beta^\tau \left(\frac{c_{t+\tau}}{c_t} \right)^{-\gamma}. \quad (6)$$

An important treatment of long-horizon returns with time-separable preferences can be found in Cochrane and Hansen (1992). They use the method of Hansen and Jagannathan (1991) to look at the implications of time-separable utility at horizons ranging from one quarter to five years. They find that the performance of time-separable utility actually *deteriorates* as the horizon lengthens. This failure of time-separable preferences at long horizons is largely caused by a high implied risk-free rate. Aggregate consumption is a (stochastically) growing series. In the time-separable model, agents seek to transfer some of the high future consumption to the present by borrowing. A counterfactually high risk-free rate is needed to discourage this borrowing. (Recall that net borrowing must equal zero in equilibrium.) However, in a stochastic model this effect is partially countered by the precautionary motive for saving: agents might wish to insure against the possibility of consumption downturns. As the horizon lengthens, this precautionary demand becomes weaker, since the empirical probability of a consumption downturn is smaller for the longer horizons. (For example, Cochrane and Hansen (1992) note that there is no five-year period in post-war US data over which aggregate consumption declines.) As a result, the equilibrium risk-free rate implied by the time-separable model is counterfactually high for longer-horizon data.

What is needed, then, is a reason why agents would be willing to save at low interest rates, even though they know that their future consumption is likely to grow. One possible reason is that agent's within-period utility-of-consumption changes through time. Preference specifications with this property include the Constantinides (1990) habit-formation preferences and Abel (1990) "catching-up-with-the-Joneses" preferences. In these specifications, the marginal utility of a given level of consumption grows through time. Agents refrain from borrowing to increase current consumption because they know that they will need the consumption more in the future. (As Weil (1989) has pointed out, the same effect could be induced by having a

subsistence point that grows deterministically in time.) To put this another way, the problem discussed by Cochrane and Hansen (1992) can be resolved, in principle, if agents are fearful, not of a decline in consumption, but of a decline in consumption *relative to some reference point*, where the reference point itself grows at the same rate that consumption grows. To resolve the risk-free rate puzzle in this way, we will utilize the Abel (1990) and Constantinides (1990) models in our empirical tests.¹³ In the following sections, we formalize these preference specifications.

3.1.2 The Abel (1990) "Catching-Up-With-the-Joneses" Preferences

Let \bar{c}_t denote the per-capita consumption at date t . The agent solves

$$\max_{\{c_{t+j}\}_{j=0}^{\infty}} U \equiv E_t \sum_{j=0}^{\infty} \beta^j \frac{(c_{t+j} - h_{t+j})^{1-\gamma}}{1-\gamma} \quad (7)$$

subject to the usual budget constraint, where

$$h_t \equiv \frac{\eta(1-\delta)}{\delta - \delta^{m+1}} \sum_{i=1}^m \delta^i \bar{c}_{t-i}, \quad \eta > 0, \quad 0 < \delta < 1. \quad (8)$$

The interpretation is that agents compare their consumption to the consumption of their neighbors (the "Joneses") in the recent past. In the formal model, the neighbors' consumption is represented by \bar{c}_t , and agents behave as if they have a subsistence point equal to η times a weighted average of the per-capita consumption levels over the past m periods. Notice that agents treat h_t as exogenous: the marginal utility of a fixed level of consumption inherits the upward trend in \bar{c}_t , but agents cannot alter the h_t process by their own actions. (Of course, in equilibrium $c_t = \bar{c}_t$.)

With Abel preferences,

$$m_{t+\tau}^{\tau} = \beta^{\tau} \frac{(c_{t+\tau} - h_{t+\tau})^{-\gamma}}{(c_t - h_t)^{-\gamma}}. \quad (9)$$

Let the value function $V(W, h)$ be defined as the maximum value of the objective function that can be attained given initial wealth W and an initial subsistence point h . We define the coefficient of relative risk aversion (denoted RRA_t) by

$$RRA_t \equiv -W \frac{V_{WW}(W, h)}{V_W(W, h)}. \quad (10)$$

¹³We note that we have also investigated preferences in which utility is based on consumption relative to a deterministic trend with similar results for the unconditional moment tests.

For Abel preferences,

$$RRA_t = \gamma \frac{c_t}{c_t - h_t} \quad (11)$$

Notice that the coefficient of relative risk aversion is time-varying, and everywhere exceeds γ . The risk-free rate puzzles will be partially resolved by the Abel specification, since, empirically, declines in $(c_t - h_t)$ are observed more frequently than declines in c_t .

3.1.3 Constantinides (1990) Habit-Formation Preferences

Constantinides (1990) models agents as maximizing an objective function of the same form as (7) with the following alternative specification for h_t :

$$h_t \equiv \frac{\eta(1-\delta)}{\delta - \delta^{m+1}} \sum_{i=1}^m \delta^i c_{t-i}, \quad \eta > 0, \quad 0 < \delta < 1. \quad (12)$$

The difference between (8) and (12) is that in (12) the stochastic subsistence point h_t is a function of the agent's own consumption c_t , rather than the per-capita consumption. The marginal rate of substitution is now

$$m_{t+\tau}^\tau = \beta^\tau \frac{MU_{t+\tau}}{MU_t} \quad (13)$$

where the marginal utility of consumption MU_t is defined by

$$MU_t \equiv (c_t - h_t)^{-\gamma} - \frac{\eta(1-\delta)}{\delta - \delta^{m+1}} \sum_{i=1}^m (\beta\delta)^i E_t [c_{t+i} - h_{t+i}]^{-\gamma} \quad (14)$$

With habit-formation preferences, agents consider the effect of their current consumption on future values of h_t . The presence of conditional expectations in equation (14) reflects this fact. These conditional expectations must be computed when we construct m_t^τ . We do this as follows. First, define the variable D_t by:

$$D_t \equiv 1 - \frac{\eta(1-\delta)}{\delta - \delta^{m+1}} \frac{\sum_{i=1}^m (\beta\delta)^i [c_{t-m+i} - h_{t-m+i}]^{-\gamma}}{(c_{t-m} - h_{t-m})^{-\gamma}}. \quad (15)$$

The variable D_t behaves as a stationary stochastic process. Equations (13) and (15) imply that, in the Constantinides model,

$$m_{t+\tau}^\tau = \beta^\tau \frac{(c_{t+\tau} - h_{t+\tau})^{-\gamma} E_{t+\tau} D_{t+\tau+m}}{(c_t - h_t)^{-\gamma} E_t D_{t+m}} \quad (16)$$

Since D_t is stationary, we can fit an autoregressive time-series model for this variable, and use the fitted values as our estimate of $E_t D_{t+m}$. For most models, the likelihood ratio statistics testing n lags against $n - 1$ lags in the autoregression for D_t (for n between 1 and 5) favor four lags. We estimate a fourth-order autoregression in D_t , and project the fitted regression m periods into the future.

We consider Constantinides preferences separately from Abel preferences for two reasons. First, it is possible that the behavior of $m_t^{\bar{\gamma}}$ in (13) may differ substantially from that implied by (9). A second, and more important reason, is that habit formation preferences do not accentuate risk aversion in the way that Abel's preferences do:¹⁴ the coefficient of relative risk aversion implied by habit-formation for a given specification of $\{\gamma, \eta, \delta, \beta\}$ is smaller than that given in (11). From the perspective of the individual agent, Abel preferences are time-separable, since a change in an individual's c_t does not affect his marginal utility with respect to c_{t+i} for $i \neq 0$. For such preferences, the coefficient of relative risk aversion equals the curvature of the per-period utility function U , as measured by $-C \frac{U''(C)}{U'(C)}$.

In contrast, habit formation preferences incorporate true time nonseparability: The marginal utility at date t is affected by changes in the state-contingent consumption plan for dates $t + i$, $i = 1, \dots, m$. In response to a wealth shock at date t , the agent with habit-formation preferences adjusts her state-contingent plans for future consumption so as to optimally adjust h_{t+i} , $i = 1, \dots, m$. This attenuates the impact of a given wealth shock on the objective function, as compared to the Abel specification. In particular, higher η does not increase RRA_t as much as in (11). While RRA_t cannot be computed analytically for our model of habit-formation, Constantinides (1990) and Ferson and Constantinides (1991) obtain closed-form solutions for RRA_t in the context of a simpler model. They show that, for preference parameters similar to ours, the mean coefficient of relative risk aversion is not too far above γ .¹⁵

¹⁴This point is extensively discussed by Constantinides (1990), Ferson and Constantinides (1991), and Boldrin, Christiano, and Fisher (1995).

¹⁵Boldrin, Christiano, and Fisher (1995) report a similar result.

3.2 Tests of Long-Horizon Pricing Using Unconditional Moments

In Section 3.3 we redo the long-horizon Hansen-Jagannathan analysis of Cochrane and Hansen (1992) using Abel (1990) and Constantinides (1990) preferences, and find that there is indeed considerable support for consumption based asset pricing models at long horizons, and very little support at short horizons. In this section we describe the long-horizon Hansen-Jagannathan tests and test-statistics we utilize.

3.2.1 The Hansen-Jagannathan Analysis

Cochrane and Hansen (1992) derive the Hansen and Jagannathan (1991) mean-variance bounds using a linear projection of a candidate stochastic discount factor on the space of portfolio payoffs. In this section, we describe an alternative derivation that is helpful in pointing out the importance of the unconditional correlation between the pricing kernel and the mean-variance efficient portfolio. In vector form, Equation (1) is:

$$E[m_t^\tau \mathbf{R}_t^\tau] = \mathbf{1}. \quad (17)$$

where \mathbf{R}_t^τ denotes the vector consisting of the gross-returns between period t and period $t + \tau$ on a set of n assets, and $\mathbf{1}$ is an n -vector of ones. Equation (17) holds for all τ , so, in the remainder of this section, we suppress the superscript τ . Equation (17) implies that:

$$\begin{aligned} \text{cov}(m, \mathbf{R}) &= E[m\mathbf{R}] - E[m]E[\mathbf{R}] \\ &= \mathbf{1} - E[m]E[\mathbf{R}] \end{aligned} \quad (18)$$

Rearranging, we obtain for any return r :

$$E[r] - E[m]^{-1} = \frac{\text{cov}(m, r)}{E[m]} = -\rho_{m,r} \left(\frac{\sigma_m \sigma_r}{E[m]} \right). \quad (19)$$

where σ_r and σ_m are the standard deviations of r and m , respectively, and $\rho_{m,r}$ is the correlation coefficient between r and m . Since the correlation coefficient $-\rho_{m,r}$ cannot be less than -1 , we see that:

$$\left(\frac{E[r] - E[m]^{-1}}{\sigma_r} \right) \leq \frac{\sigma_m}{E[m]} \quad (20)$$

where the " \leq " must hold for every asset and every feasible portfolio. Let us define the left-hand side of equation (20) as the asset's *Sharpe ratio*.¹⁶ Hansen and Jagannathan (1991) derive the strongest bound implied by (20): in effect, they evaluate the left-hand side of (20) at the portfolio return with the highest Sharpe ratio, which is simply the unconditional mean-variance efficient portfolio. This portfolio has a Sharpe ratio of

$$SR^* = \sqrt{(E[\mathbf{R}] - \mathbf{1} \cdot E[m]^{-1})' \boldsymbol{\Omega}^{-1} (E[\mathbf{R}] - \mathbf{1} \cdot E[m]^{-1})} \leq \frac{\sigma_m}{E[m]}$$

where $\boldsymbol{\Omega}$ is the covariance matrix of returns. For any given value of $E[m]$, this gives a bound on σ_m which is dependent only on the asset return first and second moments, and is independent of preferences.

Equation (19) also tells us something about the importance of the correlation between asset returns and the pricing kernel. This equation can be rewritten:

$$\frac{\sigma_m}{E[m]} = -\frac{SR_j}{\rho_{m,R_j}}$$

where SR_j denotes the Sharpe ratio of asset j . This equation says that, for any asset or portfolio of assets with $SR_j > 0$, the lower bound on $\sigma_m \rightarrow \infty$ as $\rho \uparrow 0^-$: if the pricing operator m has an arbitrarily small negative correlation with some asset with a positive Sharpe ratio, this implies an *arbitrarily* tight bound on σ_m . This clearly ties in with our results in Section 2: the positive correlation between equity returns and consumption growth suggests a negative correlation between equity returns and the marginal rate of substitution in consumption that serves as the consumption-based proxy for m . If the correlation between short-horizon equity returns and short-horizon consumption growth is positive but small, the implied bound on σ_m is likely to be much tighter than the standard Hansen-Jagannathan bound.

3.2.2 Projected Hansen-Jagannathan Bounds Tests

Following Cochrane and Hansen (1992), we note that if m satisfies (17), then so will the projection (denoted m^*) of m onto the set of asset returns and a constant:

$$m^* = a + \mathbf{R}' \cdot \mathbf{b}, \tag{21}$$

¹⁶The risk-free rate equals $E_t[m_{t+1}]^{-1}$, so our definition of the Sharpe ratio differs from the usual definition by a "Jensen's inequality" correction.

If the risk-free rate were constant, the two definitions would coincide.

where

$$m = a + \mathbf{R}' \cdot \mathbf{b} + \epsilon,$$

and $E[\epsilon\mathbf{R}] = \mathbf{0}$. If the correlation of m with the elements of \mathbf{R} is very small, then the variance of m^* will be very small relative to that of m . As in equation (20), this means that a higher value of σ_m will be required to satisfy the HJ bounds when implemented using m^* . This test therefore provides a tighter restriction than the standard HJ test. We present Hansen-Jagannathan tests both for m and m^* . We also use the procedure described in section 3.1 of Hansen, Heaton, and Luttmer (1995) to test formally whether the Hansen-Jagannathan mean-variance restrictions are satisfied.

3.2.3 A χ^2 test of the moment restriction $E[m\mathbf{R}] = \mathbf{1}$

Because we are not estimating model parameters in our test, it is straightforward to directly test the moment restriction $E[m\mathbf{R} - \mathbf{1}] = \mathbf{0}$ for the n assets being considered. If we define $\mathbf{g}_t \equiv (m_t\mathbf{R}_t - \mathbf{1})$, the sample moment estimator is $\mathbf{g}_T = (1/T) \sum_T \mathbf{g}_t$. We can construct a test statistic

$$J_T = T\mathbf{g}_T'\mathbf{S}_0^{-1}\mathbf{g}_T, \tag{22}$$

where \mathbf{S}_0 is a consistent estimator of the spectral density matrix of \mathbf{g}_t at frequency zero, and where this spectral density matrix is estimated using the Newey-West procedure.¹⁷ As shown in Hansen (1982), J_T is asymptotically distributed χ^2 with n degrees of freedom. In our case, we use two distinct asset returns, so $n = 2$.

For the time-separable model, Cochrane and Hansen show that this χ^2 test (for these two assets) strongly rejects the time-separable power utility model for γ 's of less than 50.¹⁸

3.2.4 The Hansen-Jagannathan (1997) Specification Test

As pointed out in Hansen and Jagannathan (1997), χ^2 tests of the moment restriction (17), like that discussed in Section 3.2.3, can be misleading in comparing specifications. If a model

¹⁷The number of lags used equals 8, 11, 15, and 19 lags for the quarterly, yearly, two-year, and three-year horizons, respectively.

¹⁸See Cochrane and Hansen (1992), Table 2, p. 129.

produces a highly volatile m , the eigenvalues of the spectral-density matrix may be huge, so the weighting matrix \mathbf{S}_0^{-1} in equation (22) may be close to singular. This would imply a “small” χ^2 statistic, even when the pricing errors are large.¹⁹

To help to assess whether observed low χ^2 test statistics result from a superior fit of a given model, we utilize a distance measure derived by Hansen and Jagannathan (1997).²⁰ We denote this measure the “HJ97 statistic”. This test constructs an estimator the mean squared distance between the candidate stochastic discount factor m_t and the set (\mathcal{M}) of valid pricing kernels (that is, the set of y 's that satisfy $E(y\mathbf{R} - \mathbf{1}) = \mathbf{0}$). Formally, this distance measure, $\hat{\delta}^2$, is defined by

$$\hat{\delta}^2 \equiv \min_{y \in \mathcal{M}} E(y_t - m_t)^2 \quad (23)$$

As stated, this minimization is over a set of random variables \mathcal{M} . Equation (8) in Hansen, Heaton, and Luttmer (1995) show that δ can be computed more simply as the solution to the conjugate maximization problem. We follow their procedure in computing \hat{d}_T^2 (the finite-sample analogue to $\hat{\delta}^2$) for the various models of m_t we study.²¹

HJ97 statistics are not directly comparable across different investment horizons without some normalization; here we derive a modified HJ97 statistic with appropriate normalization to allow such a comparison. The τ period long horizon misspecification measure (in our notation) is:

$$\hat{\delta}_\tau^2 \equiv \min_{y^\tau \in \mathcal{M}^\tau} E(y_t^\tau - m_t^\tau)^2 = \mu_\tau^2 + \sigma_\tau^2 \quad (24)$$

where $\mu_\tau \equiv E(\hat{y}_t^\tau - m_t^\tau)$ and $\sigma_\tau^2 = \text{var}(\hat{y}_t^\tau - m_t^\tau)$, where \hat{y} is the *argmin* of (24). If we assume that the “best-fit” pricing kernel \hat{y}^τ will not change with horizon, then the τ period pricing kernels

¹⁹For example, the χ^2 test discussed in 3.2.3 (and presented by Cochrane and Hansen (1992)), strongly rejects for small γ 's ($J_T = 5.1$ for $\gamma = 40$ for our sample), but fails to reject the time-separable model for a γ 's over 200 (J_T is equal to 0.88 for $\gamma = 240$). However, the HJ97 statistic $\hat{\delta}^2$ (defined in equation (23)) is equal to 0.257 for $\gamma = 0$ and 0.416 for $\gamma = 240$. In other words, despite the low χ^2 statistic, the m for the $\gamma = 240$ model is *further* from \mathcal{M} than the naive (and obviously sub-optimal) specification $m_t = 1$, for all t .

²⁰See also Hansen, Heaton, and Luttmer (1995) (HHL), Section 1.2).

²¹Hansen, Heaton, and Luttmer (1995) show that if $m_t \in \mathcal{M}$ (that is, if $\delta = 0$), then the limiting distribution of $\sqrt{T}\hat{d}_T$ is degenerate. For this reason, we do not use the HJ97 statistic to test whether m_t is a valid pricing kernel.

in (24) are simply products of the single period pricing kernels: $\hat{y}_t^\tau = \prod_{i=1}^\tau \hat{y}_{t+i-1}^1$ and $m_t^\tau = \prod_{i=1}^\tau m_{t+i-1}^1$. Since, for small τ , both y and m should be close to one, and since:

$$\prod_{i=1}^N (1 + \delta_i) \approx 1 + \sum_{i=1}^N \delta_i$$

we have that:

$$\hat{y}_t^\tau - m_t^\tau = \prod_{i=1}^\tau \hat{y}_{t+i-1}^1 - \prod_{i=1}^\tau m_{t+i-1}^1 \approx \sum_{i=1}^\tau (\hat{y}_{t+i-1}^1 - m_{t+i-1}^1)$$

Thus, under this approximation,

$$\mu_\tau \approx \tau \mu_1,$$

and

$$\sigma_\tau^2 = \tau \cdot \sigma_1^2 \cdot VR(\tau),$$

where the variance ratio $VR(\tau)$ is defined in the usual way as the ratio of the variance of $(y^\tau - m^\tau)$ to the variance of $(y^1 - m^1)$, divided by τ :

$$VR(\tau) \equiv \frac{\sigma_\tau^2}{\tau \cdot \sigma_1^2}.$$

It follows that the HJ97 statistic can be written

$$\hat{\delta}_\tau^2 = \mu_\tau^2 + \sigma_\tau^2 \approx \tau^2 \cdot \mu_1^2 + \tau \cdot \sigma_1^2 \cdot VR(\tau). \quad (25)$$

We now define the *modified HJ97 statistic* $\hat{\delta}_\tau^\dagger$ as

$$(\hat{\delta}_\tau^\dagger)^2 \equiv \frac{\mu_\tau^2}{\tau^2} + \frac{\sigma_\tau^2}{\tau} = \mu + \sigma^2 \cdot VR(\tau) \quad (26)$$

For $\tau = 1$, $(\hat{\delta}_\tau^\dagger)^2$ is equal to the standard HJ97 statistic in (24). To implement (26), we compute \hat{y} as the *argmin* of the sample version of equation (24).²², we calculate the sample $\hat{\mu}^\tau$ and $\hat{\sigma}^\tau$, and form the square-root of the sample version of our modified statistic, \hat{d}_τ^\dagger . In the empirical section, we use the modified HJ97 statistic \hat{d}_τ^\dagger to assess the fit of the model at different horizons.

The modified HJ97 statistic tells us exactly what we want to know about how the characteristics of $(y - m)$ change as the test horizon lengthens: if $(y_t - m_t)$ is characterized by

²²HHL show that the $\hat{y} \in \hat{\mathcal{M}}$ which minimizes the sample version of equation (24) is $\hat{y}_t = m_t - \mathbf{R}'\boldsymbol{\alpha}$, where $\boldsymbol{\alpha}$ is the vector of Lagrange multipliers of the original minimization problem, and where the F.O.C. defining $\boldsymbol{\alpha}$ is given in their equation (10). Therefore, $(\hat{y}_t - m_t) = \hat{\boldsymbol{\alpha}}'\mathbf{R}$. μ_τ and σ_τ are the sample mean and standard deviation of this series.

rapid, negatively autocorrelated movements which are “washed-out” at longer horizons, then the variance-ratio in (25) will be small and \hat{d}_τ^\dagger will be small for long horizons. This would tell us that eliminating high frequency components improves the fit of the model. If, on the other hand, $(y_t - m_t)$ is characterized by positively correlated, long-swings, then the variance-ratio in (25), and \hat{d}_τ^\dagger will be large for long horizons, and the modified HJ97 statistic will indicate a poorer fit to the data at longer-horizons.

3.3 Empirical Results

We now present the results from our empirical analysis of the models presented in Section 3.1, using the tests presented in Section 3.2. We use quarterly data from 1952-1997. Detailed description of the data can be found in the Data Appendix. We follow the standard practice in the literature in using consumer expenditures on nondurables and services (“ND&S”) as our primary measure of consumption. We also perform the analysis for consumer expenditures on nondurables (“ND”). Finally, The data compiled by the Bureau of Economic Analysis data for consumer expenditures on services includes a series measuring the imputed rental value of owner-occupied housing. Rather than being compiled from observed surveys on consumer expenditure, this series is a construct. It has rather different properties from other consumption series. In particular, Marshall and Parekh (1998) note that the covariance between the growth rate of this series and equity returns is negative. To allow for the possibility that this series is substantially mis-measured, our third measure of consumption is ND&S minus this imputed rental value of owner-occupied housing.


3.3.1 Time-Separable Preferences

We examine time-separable preferences for four horizons: one quarter, one year, two years, and three years. We set β equal to 1.0, and we vary the curvature parameter γ . The behavior of this model for ND&S consumption is displayed in Figure 5.²³ The solid “U-shaped” curves in Figure 5 plot the Hansen-Jagannathan bounds, calculated using two returns: the real CRSP

²³We also conducted this analysis using non-durable consumption and using ND&S minus the service flow from owner-occupied housing. The results are very similar to those for ND&S, and are not displayed.

Figure 5: **Hansen and Jagannathan bounds tests for Time-Separable Preferences, Non-Durable and Services Consumption**

The four solid curves in both panels are the Hansen-Jagannathan bounds on the quarterly standard deviation of the pricing kernel, inferred from the asset return data at the four horizons: one quarter, one year, two years, and three years. In the top panel, the dotted lines plot the means/standard deviations for the time-separable model marginal rate of substitution for the following four horizons: one quarter (“+”), one year (“*”), two years (“circle”), and three years (“x”). We set $\beta = 1$ and we let the value of γ range from 0 to 200; for each line the spacing between the symbols is $\gamma = 5$. In the bottom panel, the dashed lines are generated by running a regression of the candidate discount factors on the set of real returns (in this case the VW index and the T-Bill rate). The means and standard deviations of the fitted regressions are plotted for $\gamma = 0, \dots, 200$. Consumption is measured as expenditures on consumer nondurables plus services; the time period is 1947:1-1997:4.



VW index return and the real 3-month T-Bill return. To facilitate comparison across different investment horizons, we normalize the plots to put both mean and variance of the pricing kernel in quarterly terms. To do this, we divide the resulting lower bound on σ_m by $\sqrt{\tau}$, where τ is the horizon length in quarters (1, 4, 8 or 12 in this plot). The mean is normalized by plotting $E[m]^{1/\tau}$. In the upper panel of Figure 5 we plot the mean-standard deviation loci of m_t implied by time-separable preferences with ND&S consumption, for the four horizons and for values of γ ranging from 0 to 200. In the lower panel, we project the time-separable model's pricing kernel onto the set of asset returns as described in Section 3.2.2, and plot the mean and standard deviation of m^* (defined in equation (21) against the same H-J bounds.

Figure 5 confirms the failure of the time-separable model that has been noted by many previous studies. At every horizons from one quarter through three years, the (mean, standard deviation) loci of the pricing kernel implied by this model are substantially outside of the Hansen-Jagannathan bound. As in Cochrane and Hansen (1992), the longer horizons are further from the bound than the quarterly horizon. We test formally these Hansen-Jagannathan mean-variance restrictions using the procedure described in Hansen, Heaton, and Luttmer (1995).²⁴ The failure of the time-separable model is generally confirmed. In particular, the Hansen-Jagannathan restrictions for the projected kernel are rejected at the 1% marginal significance level for all horizons except when γ is near zero²⁵ or, for the quarterly horizon, for values of γ exceeding 200.²⁶

One should not, however, take these results as indicating that time-separable preferences can fit the data with γ 's near zero or (for the quarterly horizon) near 200. When χ^2 statistics are calculated using equation (22). The statistics reject the model for all values of γ . In particular, the smallest values of these χ^2 statistics are 24.8, 16.8, 14.9, and 12.32 for horizons of one quarter, one year, two years, and three years, respectively. All of these values exceed 10.60, which is the 0.5% critical value for a χ^2 statistic with two degrees of freedom. The modified HJ97 statistics, given in equation (26), tell the same story. These statistics range between 0.273 and 0.437, implying that the root-mean-squared deviation of the model's pricing

²⁴See Hansen, Heaton, and Luttmer (1995), section 3.1.

²⁵When $\gamma = 0$ (risk-neutrality), this test has p-values of 0.073, 0.048, 0.074, and 0.056 for horizons of one quarter, one year, two years, and three years, respectively.

²⁶The p-value for the quarterly horizon when $\gamma = 200$ is 0.087.

kernel from the set of valid kernels is between 2730 bp and 4370 bp per quarter.

3.3.2 Time-Nonseparable Preferences: An Overview of the Empirical Results

In this section, we conduct an analysis of the Abel and Constantinides models analogous to that done in section 3.3.1 for time-separable preferences. That is, we look at the models' performance as curvature increases. While the model with time-separable preferences has only two parameters, β (the subjective discount rate) and γ (the curvature parameter), the Abel and Constantinides models have five parameters: β , γ , δ (the rate at which the habit stock depreciates), η (the ratio of mean consumption to mean habit stock) and m (the number of lagged consumptions that enter the habit stock). In this subsection, we proceed somewhat informally: We simply fix $\{\beta, \delta, \eta, m\}$ at values where the models perform fairly well²⁷ and we vary γ , as in figure 5. In section 4, below, we consider the models' performance when the parameters are chosen optimally.

Note that equation (14) for the marginal utility of consumption in the Constantinides model only makes sense if this marginal utility is positive. If the second term in equation (14) exceeds the first term, the expected effect of additional consumption on the future habit stocks would outweigh the effect on current period utility. In that case, it would be optimal for the consumer to reduce consumption via free disposal. In addition, the right-hand side of equations (9) or (14) could be negative or complex if $(c_t - h_t)$ is negative. We discard any parameterizations of these models where either of these anomalies occurred.

3.3.3 The Abel Model

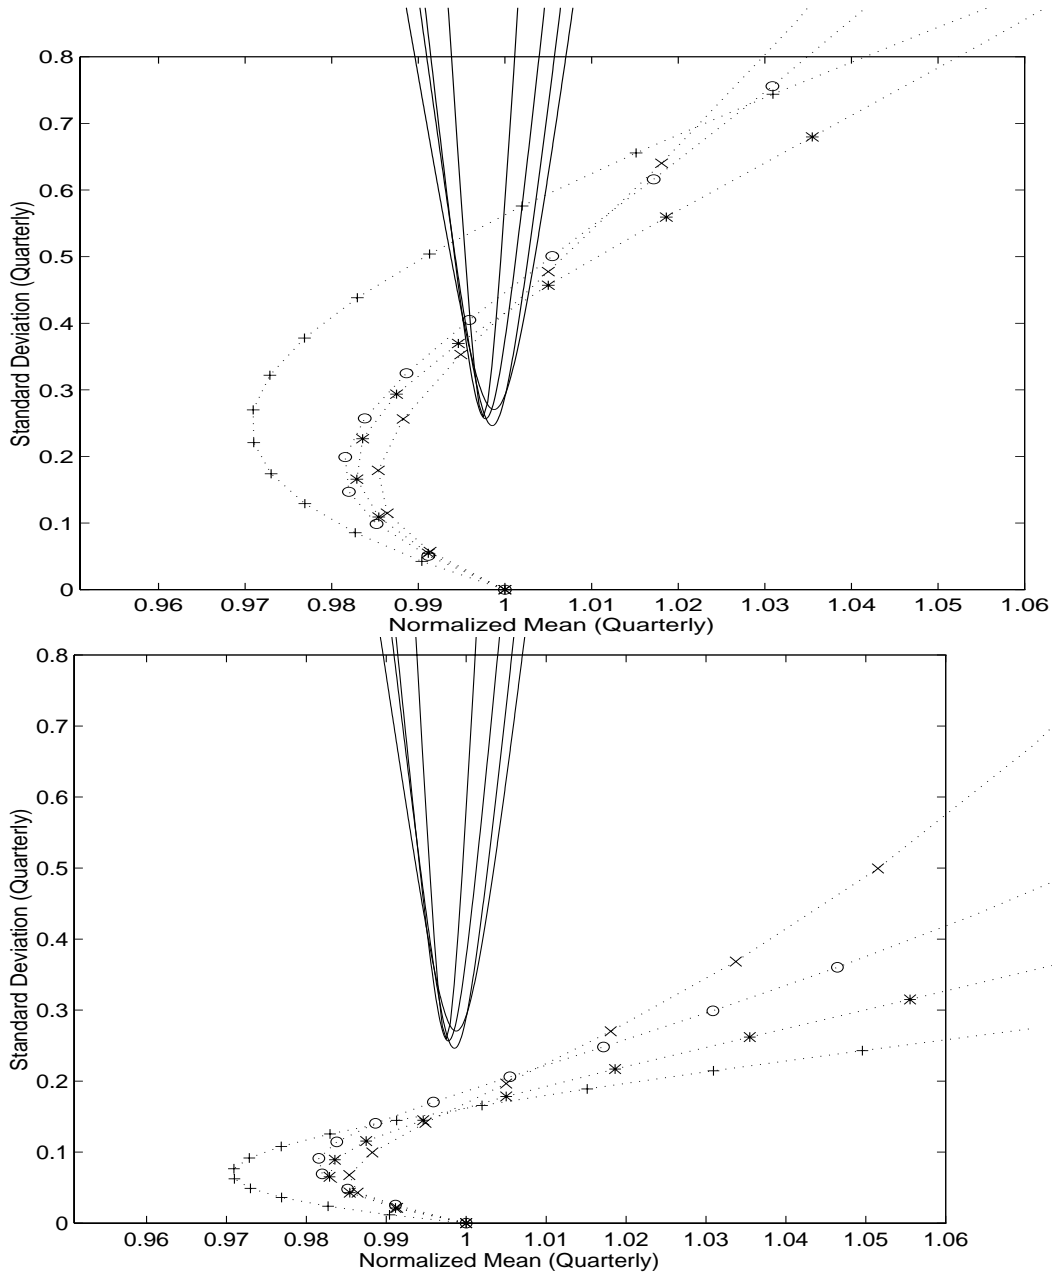
For the Abel model, we set $\eta = 0.8$; $\delta = 0.9$; $\beta = 1$; and $m = 20$. Figure 6 is analogous to Figure 5. The upper panel plots the mean-standard deviation loci of m_t implied by Abel with ND&S consumption, for the four horizons, and for various values of the curvature parameter γ . The lower panel is the analogous plot for the projection of Abel model pricing kernel onto the set of asset returns.

The upper panel shows that the model pricing kernels with ND&S consumption satisfy the

²⁷We conduct a preliminary search in which we set $\beta = 1.0$ and we vary η and δ from 0.5 through 0.9, we vary γ from 0 through 30, and we vary m from 2 through 20.

Figure 6: Hansen and Jagannathan bounds tests for Abel Preferences, for Non-Durable and Services Consumption

The four solid curves in both panels are the Hansen-Jagannathan bounds on the quarterly standard deviation of the pricing kernel, inferred from the asset return data at the four horizons: one quarter, one year, two years, and three years. In the top panel, the dotted lines plot the means/standard deviations for the Abel model marginal rate of substitution, as given in equation (9), for the following four horizons: one quarter (“+”), one year (“*”), two years (“circle”), and three years (“x”). The model parameters are $\eta = 0.8$, $\delta = 0.9$, $m = 20$, $\beta = 1$. The value of γ ranges from 0 to 30; for each line the spacing between the symbols is $\gamma = 2$. In the bottom panel, the dashed lines are generated by running a regression of the candidate discount factors on the set of real returns (in this case the VW index and the T-Bill rate). The means and standard deviations of the fitted regressions are plotted for $\gamma = 0, \dots, 30$. Consumption is measured as expenditures on consumer nondurables plus services; the time period is 1947:1-1997:4.



HJ bounds at the yearly (“*”) two-year (“circle”), , and three-year (“x”) horizons when $\gamma = 10-14$. The quarterly horizon (“+”) requires a somewhat higher value of γ (over 20). The lower panel, where the pricing kernels are projected onto the space of asset returns, give a somewhat different story. While the two-year horizon comes closest to the HJ bounds, none of the point estimates implied by the model actually satisfy these bounds. The reason is that the correlation between consumption growth and returns is small. However, when the Hansen-Jagannathan mean-variance restrictions are tested formally using the Hansen, Heaton, and Luttmer (1995) procedure, the distance between the HJ bound and the (mean,variance) loci of the pricing kernels is significantly different from zero only for the quarterly horizon (and perhaps for the 3-year horizon) at low values of γ . In particular, for γ 's in the range of 6-24, p-values of above 10% are achieved for the three longer horizons. (For the quarterly horizon, γ 's greater than 10 are required.) This means that, for moderately high values of γ , none of the models are significantly outside the HJ bounds according to this test.

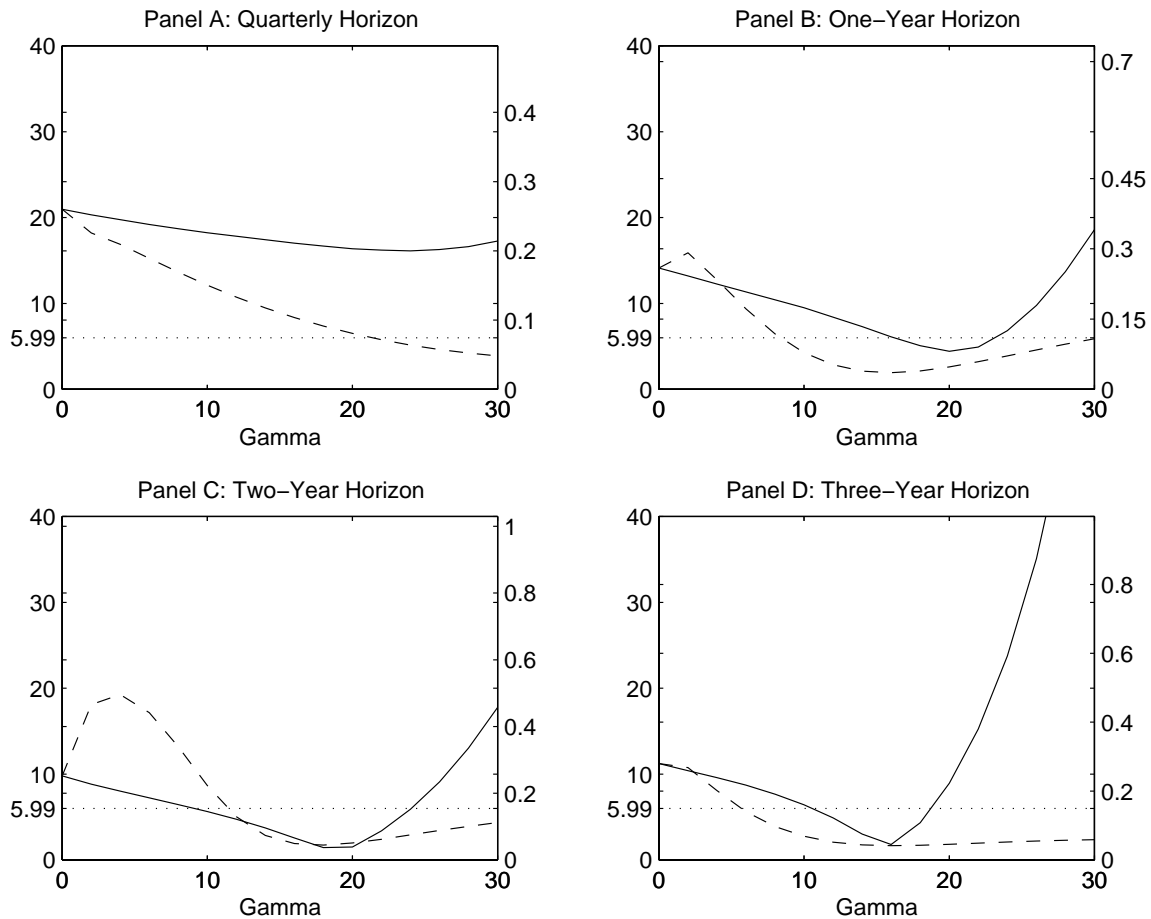
The χ^2 statistics, calculated using equation (22) give similar results. These statistics are graphed in Figure 7. For each of the four return horizons, the dashed line displays the value of the χ^2 statistic as a function of γ . The scale for this statistic is given at the left-hand side of each graph. The horizontal dotted line indicates the value of 5.99, which is the 5% critical value for a χ^2 random variable with two degrees of freedom. This critical value is given only as a point of reference. If the parameters were chosen before looking at the data, the $\chi^2(2)$ would be the appropriate distribution under the null hypothesis that the model fits in population, since equation (2) must hold for each of two assets. However, the parameters used in this section were chosen after an informal grid search, so the $\chi^2(2)$ critical values should only be used as an informal guide.

In figure 7, the value of the χ^2 statistic drops rapidly for all horizons as γ increases. For high enough values of γ , there seems to be little evidence against the model. Note in particular the rapid monotonic decline for the quarterly horizon (panel A). Results like these lead some to argue that models with time non-separabilities are consistent with the behavior of short-term asset returns.²⁸ An alternative interpretation, put forth by Hansen and Jagannathan (1997), is

²⁸See, for example, Ferson and Constantinides (1991) and Campbell and Cochrane (1994).

Figure 7: χ^2 statistics and Hansen-Jagannathan (1997) statistics for Abel Preferences.

The figure plots the χ^2 statistic (dashed line) and the HJ97 statistic (solid line) implied by the Abel model at the quarterly (panel A), one-year (panel B), two-year (panel C), and three-year (panel D) horizon for values of γ ranging from 0 to 30 (horizontal axis). The scale for the χ^2 statistic is on the left-hand side of the graph; the scale for the HJ97 statistic is on the right-hand side of the graph. The horizontal dotted line indicates the value of 5.99, which is the 5% critical value for a χ^2 random variable with two degrees of freedom. The model parameters are $\eta = 0.8$, $\delta = 0.9$, $m = 20$, $\beta = 1$. Consumption is measured as expenditures on consumer nondurables plus services; the time period is 1947:1-1997:4.



that the χ^2 statistic lacks power. In figure 7, the solid lines plot the modified HJ97 statistics as γ increases. The scale for these statistics, given at the right-hand side of each graph, is chosen so that the HJ97 statistic and the χ^2 statistic are at the same level for $\gamma = 0$. According to these results, only the longer horizons with γ around 16-20 perform particularly well. A useful point of comparison for these statistics is the value when $\gamma = 0$, which implies a constant IMRS of one. The HJ97 statistics for $\gamma = 0$ are therefore the root-mean square (RMS) distance from the set \mathcal{M} of valid pricing kernels to $m_t = 1$. The HJ97 statistics for the quarterly test of the Abel model show that the RMS distance from the model to \mathcal{M} is not much smaller than the distance to $m_t = 1$ for *any* value of γ . However, for the 1 and 2 year horizons, the statistics are about a factor of six smaller than for the model with constant m_t : this indicates that m is indeed getting close to the set \mathcal{M} . The minimal value of the modified HJ97 statistic at the two-year horizon (figure 7, panel C) is 0.037, occurring at $\gamma = 18$. That is, the minimal RMS error is on the order of 3.7%. The RMS error for the three-year horizon is approximately 4.5%, when γ equals 16.

We have conducted a similar analysis of the Abel model using the alternative consumption measures.²⁹ When consumption is measured as purchases of consumer nondurables, the HJ97 statistics tend to be less favorable to the model. Higher values of γ are required to get a reasonable fit at the two-year horizon, and the three-year horizon does poorly. In contrast, when the service flow from owner-occupied housing is omitted from the standard ND&S measure of consumption, the model tends to perform somewhat better. In particular, rather low values of the HJ97 statistic are found for the two- and three-year horizon when γ equals 14.³⁰ This is of interest because the service flow from owner-occupied housing is not an observed series, but is a construct that may be fraught with measurement error.

3.3.4 Constantinides Habit Persistence


We conclude this section by performing an analogous study of a particular parameterization of Constantinides preferences. We measure consumption as ND&S, and set $\delta = 0.9$, $\eta = 0.8$,

²⁹Detailed tabulations of these results are available from the authors.

³⁰With this measure of consumption, the HJ97 statistic when $\gamma = 14$ is 0.038 for the two-year horizon and 0.053 for the three-year horizon.

Figure 8: **Hansen-Jagannathan bounds tests for Constantinides Preferences, for Non-Durable and Services Consumption**

The description for figure 6, applies here also, except that the preferences here are Constantinides; the model parameters are $\eta = 0.8$, $\delta = 0.9$, $m = 2$, $\beta = 1$; the value of γ ranges from 0 to 18; and the spacing between the crosses is $\gamma = 1$. This figure is for Non-Durable and Services Consumption.



and $m = 2$. The Hansen-Jagannathan bounds plots for are presented in Figure 8. Unlike the Abel model displayed in Figure 6, the point estimates for this parameterization of the Constantinides model satisfy the HJ bounds only at the two-year horizon. When the pricing kernels are projected onto the space of asset returns (lower panel), none of the point estimates satisfy the HJ restrictions. As with the Abel model, the two-year horizon comes closest to the HJ bounds. However, the Hansen-Jagannathan mean-variance restrictions not rejected statistically by the Hansen, Heaton, and Luttmer (1995) test except at the three year horizon.

Figure 9 is the analogue for Constantinides preferences to figure 7. The χ^2 test appears less favorable to the quarterly horizon than with Abel preferences: In panel A of figure 9 the χ^2 statistics exceed the 5.99 critical value for all values of γ . The three-year horizon does somewhat more poorly according to this criterion with Constantinides preferences, as compared to Abel preferences. However, both the one- and two-year horizons do quite well for values of γ exceeding 4. Moreover, the modified HJ97 statistics (solid lines) indicate that the χ^2 statistics are small because the distance between m and \mathcal{M} is indeed shrinking. For example, the RMS distance between the model's IMRS and the set of valid pricing kernels is only 0.018 for the two-year horizon with γ equal to 11. As with the Abel model, the Constantinides model does less well when consumption is measured by nondurables only. When housing services are omitted from the data, the results are similar to those for ND&S.

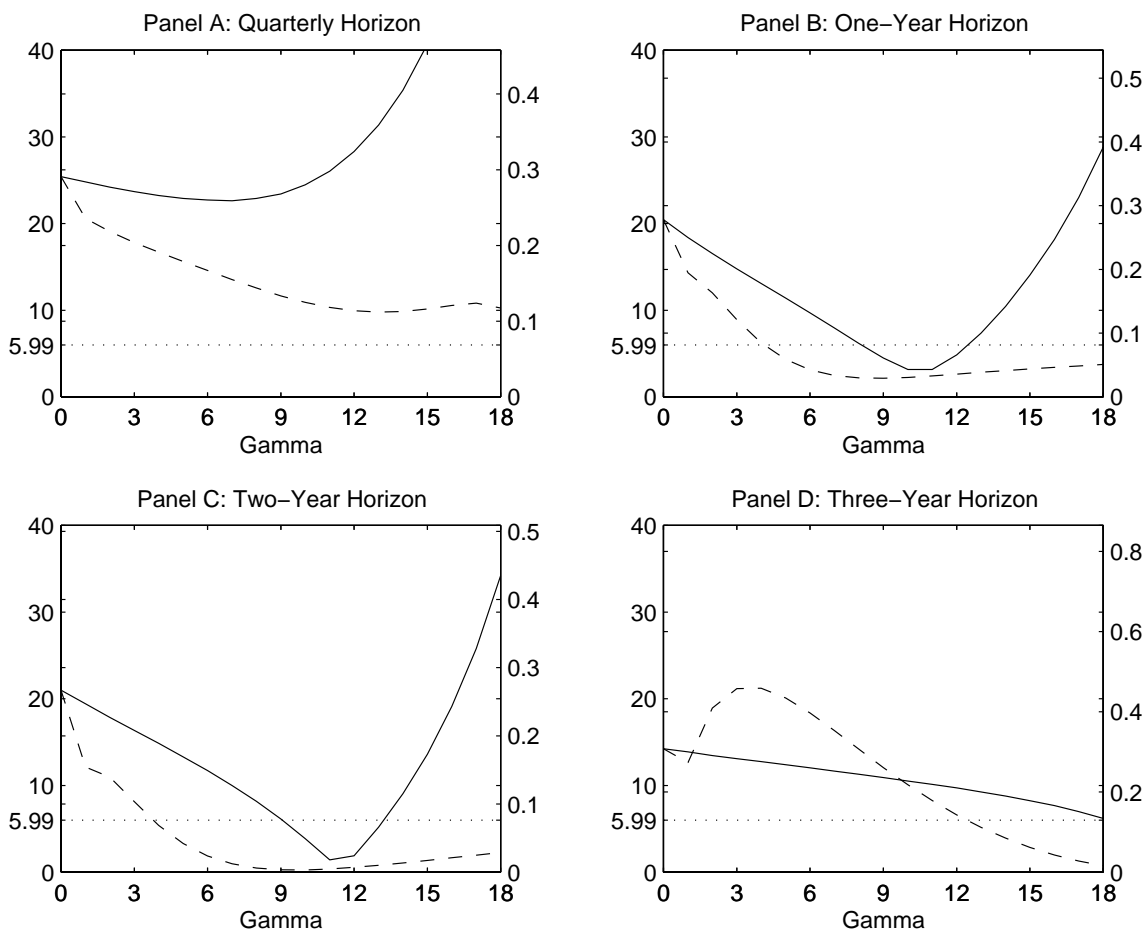
To summarize the evidence from this section, the graphical H-J analysis, the χ^2 statistics, and the HJ97 statistics suggest that the Abel model performs best at the two-year horizon and the Constantinides model does fairly well at both the one- and two-year horizons.

4 Evaluation of Time Non-Separable Preferences When Model Parameters Are Chosen Optimally

In the previous section, we simulate the Abel and Constantinides models over a grid of parameters. We found that, broadly speaking, these models appeared to perform poorly at the quarterly horizon, fairly well at the two-year horizon, and with intermediate performance at the one- and three-year horizons. In this section, we ask how well these models perform when the parameters are chosen optimally. In particular, for each model we wish to see whether

Figure 9: χ^2 statistics and Hansen-Jagannathan (1997) statistics for Constantinides Preferences.

The figure plots the χ^2 statistic (dashed line) and the HJ97 statistic (solid line) implied by the Constantinides model at the quarterly (panel A), one-year (panel B), two-year (panel C), and three-year (panel D) horizon for values of γ ranging from 0 to 30 (horizontal axis). The scale for the χ^2 statistic is on the left-hand side of the graph; the scale for the HJ97 statistic is on the right-hand side of the graph. The horizontal dotted line indicates the value of 5.99, which is the 5% critical value for a χ^2 random variable with two degrees of freedom. The model parameters are $\eta = 0.8$, $\delta = 0.9$, $m = 2$, $\beta = 1$. Consumption is measured as expenditures on consumer nondurables plus services; the time period is 1947:1-1997:4.



there are parameter configurations that set the HJ97 statistic to zero (indicating that the mean residuals from the asset pricing Euler equations equal zero in sample). If a zero value of this criterion is achieved, we wish to see whether the parameter that achieve this value are "plausible" in some intuitive sense.³¹


In conducting this exercise, we note that the model is under-determined. We have only two unconditional Euler equations (one for the 3-month T-bill return, and one for the equity return). However, in each of the two models we have five parameters to choose: $\{\beta, \gamma, \eta, \delta, m\}$. We set the quarterly subjective discount factor $\beta = .99$ (implying a yearly discount rate of approximately 4%). We treat m , the number of lagged consumption in the habit-stock formation, as an approximation to an infinite lag. We do so by letting m depend on δ : we set m as the smallest number such that $\delta^m < 0.05$. That is, we choose an m sufficiently big that the discarded lags (those greater than m) all have discount factors less than 5%. This still leaves us with three parameters $\{\gamma, \eta, \delta\}$ and only two model restrictions. Since this is a highly nonlinear model, one could find no parameter combinations that set the HJ97 statistic to zero or one could find multiple parameter combinations that achieve this goal. To explore this question we fix δ at three values $\{0.6, 0.7, 0.8\}$, and for each value we minimize the HJ97 statistic over parameters $\{\gamma, \eta\}$. When $\delta = 0.8$ (the maximum value we use), $m = 14$, implying that the habit stock incorporates three-and-a-half years of lagged consumptions.

Before describing our general results, we illustrate the behavior of a particular model. In figure 10 we consider how well the Constantinides model performs for two-year returns when $\delta = 0.6$. We consider values of γ from zero to 40 (horizontal axis), and for each value of γ we find the value of η that minimizes the HJ97 statistic. Note that both γ and η are curvature parameters. One might think that there is a trade-off between these two parameters: a high value of γ acting similarly to a high value of η . In the bottom panel of figure 10 we see that this is indeed the case: The optimal choice of η declines as γ increases. It does not then follow, however, that γ and η are not separately identified. As one can see in the top panel of figure 10, there is a unique $\{\gamma, \eta\}$ combination that minimizes the HJ97 criterion. This pattern holds

³¹This exercise is analogous to that performed by Campbell (1998) for time-separable preferences. Using a variety of data sets from different countries, he computes the risk aversion parameter needed to match the mean equity premium. His test of the model is whether the requisite level of risk aversion is plausible from an economic standpoint.

Figure 10: **Optimal Choice of Curvature Parameters for Constantinides Model, Two Year Horizon**

For a grid of γ 's ranging from zero to 40, the Hansen-Jagannathan (1997) statistic ("HJ97") is minimized by choice of parameter η . The Constantinides model of preferences is used with consumption measured by ND&S. Other model parameter are set as follows: $\beta = 0.99$; $\delta = 0.6$, $m = 6$. The top panel of this figure plots the minimized HJ97 statistic for each value of γ . The bottom panel gives the minimizing choice of η for each value of γ . The time period is 1947:1-1997:4.



generally: In all model variants and all return horizons, for each δ there is a uniquely-identified optimal $\{\gamma, \eta\}$ combination.

Table 2 gives the best-performing configuration of $\{\delta, \gamma, \eta\}$ for each model and each horizon. (When multiple values of δ achieve a zero value for the HJ97 statistic, we display the parameterization with the lowest value of γ .) The results broadly confirm the patterns discussed in the previous section. First, for no model can we find a parameter combination that performs acceptably at the quarterly horizon. In the Abel model, the minimum HJ97 statistic for quarterly returns is 0.214; for the Constantinides model, this minimum HJ97 value is 0.115. Varying the number of lags in the habit stock does not improve the performance of these models, nor does using alternative measures of consumption.³² We conclude that these models are inconsistent with the unconditional moments of short-horizon asset returns.

Second, for each model and each measure of consumption there exist parameter combinations $\{\gamma, \eta, \delta\}$ that set the HJ97 statistic to zero for the two-year horizon. For the Constantinides model, this good fit to the data does not require an excessive value of the curvature parameter γ . In particular, when $\delta = 0.8$, the Constantinides model requires γ of only 7.4. The Abel model requires somewhat higher values of γ to set the HJ97 statistic to zero. The best fit for the Abel model at the two-year horizon occurs when $\delta = 0.6$. For this value of δ , the optimal γ equals 17.9. In that sense, one can argue that the Abel model has more difficulty than the Constantinides model in fitting the unconditional moments of the asset-return data. Furthermore, in the Abel model the coefficient of relative risk aversion is time-varying, but always exceeds γ , often by a good deal. (See equation (11).) For example, the mean coefficient of relative risk aversion implied by the parameters that achieve a zero HJ97 statistic for the two-year horizon ($\delta = 0.6, \gamma = 17.9, \eta = 0.85$) is 114. Thus, the best-fitting parameterizations of this model do imply extremely high (perhaps implausible) risk aversion. In contrast, it has been noted by Constantinides (1990), Ferson and Constantinides (1991), and Boldrin, Christiano, and Fisher (1995) that habit formation in the Constantinides model does not accentuate risk aversion to the extent that it does in the Abel model. While a closed-form expression for

³²The minimal HJ97 statistics obtained for the quarterly horizon using nondurable consumption only are 0.232 and 0.136 for the Abel and Constantinides models, respectively. When housing services are omitted from ND&S, the corresponding statistics are 0.215 and 0.117.

the coefficient of relative risk aversion is not available for the Constantinides model, numerical analyses of simpler versions of the model by these authors suggest that the mean coefficient of relative risk aversion is only slightly higher than γ . We conclude that the Constantinides model can match the unconditional properties of these asset returns at the two-year horizon without imposing extreme levels of risk aversion.

Third, in both models we can find parameter combinations $\{\gamma, \eta, \delta\}$ that set the HJ97 criterion to zero for the one-year horizon. This generally holds for other measures of consumption as well.³³ However, the requisite value of γ needed at this one-year horizon is always substantially higher than that needed for the two-year horizon. Specifically, the best-performing parameter configuration for the Abel model at the one-year horizon sets γ equal to 72. As with the two-year horizon, the Constantinides model requires a substantially lower value of 17.2. Still, this value of γ exceeds that required to fit the data at the two-year horizon. Fourth, the Constantinides model does not set the HJ97 statistic to zero at the three-year horizon for any value of δ . The Abel model can achieve a zero value for the HJ97 statistic at the three-year horizon (when δ equals 0.7 or 0.8), if a high value of γ (in excess of 68) is assumed.

Finally, the results do seem sensitive to the number of lagged consumptions (" m ") used to construct the habit stock. While the Constantinides model performs fairly well at the two-year horizon when $\delta = 0.6$ or 0.8, it does rather poorly when $\delta = 0.7$. The problem is not with the value of δ *per se*, but with the value of $m = 9$ implied by $\delta = 0.7$. When we fix $\delta = 0.7$ but set $m = 6$, the model behaves much like the case in Table 2 with $\delta = 0.6$. Similarly, when $m = 10$ (but δ remains fixed at 0.7), the model resembles the case in Table 2 with $\delta = 0.8$.

5 Conclusions

In this paper, we ask whether consumption-based pricing models work better at longer horizons than at the quarterly horizon. Our motivation is that if frictions, transactions costs, or durability affect the comovements of consumption growth and asset returns, they should primarily affect the higher frequency components.

³³An exception is the Abel model using nondurable consumption, where the minimal value of the HJ97 statistic is 0.015 (when δ is set to 0.8).

Table 2: **Minimum HJ97 statistics for Abel and Constantinides Models**

For each model, δ is fixed at 0.6, 0.7, or 0.8, m is set to the minimum value such that $\delta^m < 0.05$, and the HJ97 statistic is minimized over $\{\gamma, \eta\}$. These tables report the minimal value of the test statistic obtained, for horizons of one quarter and 1, 2, and 3 years. Consumption is measured as ND&S. When the minimized HJ97 statistic equals zero for more than one value of δ , the case with the lower value of γ is displayed.

Horizon (years)	Abel Model					Constantinides Model				
	HJ _{min}	δ	γ	η	m	HJ _{min}	δ	γ	η	m
0.25	0.214	0.8	23.0	0.79	14	0.115	0.8	28.9	0.77	14
1	0.015	0.8	72.1	0.59	14	0.000	0.8	17.2	0.80	14
2	0.000	0.6	17.9	0.85	6	0.000	0.8	7.4	0.87	14
3	0.000	0.6	68.2	0.79	6	0.014	0.6	58.8	0.85	6

We first show that consumption-growth and equity returns are virtually uncorrelated at high frequencies. However, we find that at lower frequencies, corresponding to swings longer than one and one-half years, the two series are highly correlated. We then test three models of the pricing kernel: time-separable power utility; the Abel (1990) “Catching up with the Joneses” preferences; and the Constantinides (1990) habit-formation preferences. We find while all models perform poorly at quarterly horizons, the Abel (1990) and Constantinides (1990) models perform well at longer-horizons. The Hansen and Jagannathan (1991) bounds, modified to take account of correlation the correlation between the pricing kernel and asset returns, are satisfied. A χ^2 moment restriction test is not rejected, and the modified Hansen, Heaton, and Luttmer (1995) specification test suggests that our inability to reject the model is not due high volatility of the pricing kernel.

Our results raise numerous questions for further research. First, it suggests that correlation puzzle may at least partially be due to frictions that disrupt the high-frequency co-movements of marginal utility growth and returns. However, it also necessary to first understand why past attempts at modeling frictions to explain the equity-premium puzzle have not been entirely successful.

The Abel and Constantinides models motivate agents to save at relatively low interest rates, even though consumption grows, because the marginal utility-value of a given level of consumption also grows. There may be other ways to motivate saving, such as by looking

carefully at life-cycle models of saving and investment.

Finally, we consistently find that the models with time-nonseparabilities perform better at the two-year horizon than at the three-year horizon. (The Constantinides model does rather poorly at the three-year horizon.) This represents somewhat of a puzzle. If the only reason for the poor performance of consumption-based models at short horizons is transient, short-term frictions, then one would expect the models' performance to improve monotonically as the horizon lengthens. One possible reason is that consumers are particularly averse to risk at business cycle frequencies. At the three year horizon, much of the business cycle behavior of consumption growth and asset returns is filtered out. However, it is unclear why consumers would be relatively insensitive to lower-frequency risk. Whether this reflects a degree of myopia or a hitherto un-modeled aspect of preferences remains to be determined.

References

- Abel, Andrew B., 1990, Asset prices under habit formation and catching up with the Joneses, *American Economic Review* 80, 38–42.
- Baxter, Marianne, 1994, Real exchange rates and real interest rate differentials: Have we missed the business cycle relationship?, *Journal of Monetary Economics* 33, 5–37.
- , and Robert King, 1994, Measuring business cycles: Approximate band-pass filters for economic time series, unpublished manuscript, Department of Economics, University of Virginia.
- Bloomfield, Peter, 1976, *Fourier Analysis of Time Series: An Introduction* (John Wiley & Sons: New York).
- Boldrin, Michele, Lawrence Christiano, and Jonas Fisher, 1995, Asset pricing lessons for modeling business cycles, unpublished manuscript.
- Brainard, W.C., M.D. Shapiro, and J. Shoven, 1991, Market value and fundamental value, in W.C. Brainard, W.D. Nordhaus, and H.W. Watts, ed.: *Money, Macroeconomics, and Economic Policy: Essays in Honor of James Tobin* (MIT Press: Cambridge).
- Breeden, Douglas T., 1979, An intertemporal asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* 7, 265–96.
- Campbell, John, 1998, Asset prices, consumption, and the business cycle, NBER Working Paper No. 6485.
- Campbell, John Y., and John H. Cochrane, 1994, By force of habit: A consumption based explanation of aggregate stock market behavior, Working paper, University of Chicago.
- Cecchetti, Stephen G., Pok-Sang Lam, and Nelson C. Mark, 1993, The equity premium and the risk-free rate: Matching the moments, *Journal of Monetary Economics* 31, 21–46.
- Chen, Nai-Fu, 1991, Financial investment opportunities and the macroeconomy, *Journal of Finance* 46, 529–55.
- Cochrane, John H., 1989, The sensitivity of tests of the intertemporal allocation of consumption to near-rational alternatives, *American Economic Review* 79, 319–337.
- , and Lars Peter Hansen, 1992, Asset pricing explorations for macroeconomics, in Olivier Blanchard, and Stanley Fischer, ed.: *1992 NBER Macroeconomics Annual* (Cambridge: MIT Press).
- Constantinides, George M., 1990, Habit formation: A resolution of the equity premium puzzle, *Journal of Political Economy* 98, 519–543.
- , and Darrell Duffie, 1996, Asset pricing with heterogeneous consumers, *Journal of Political Economy* 104, 219–240.
- Daniel, Kent D., and David Marshall, 1997, Equity premium puzzle and risk-free rate at long horizons, *Macroeconomic Dynamics* 1, 452–484.

- Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.
- Ferson, Wayne E., and George M. Constantinides, 1991, Habit persistence and durability in aggregate consumption; empirical tests, *Journal of Financial Economics* 29, 199–240.
- Grossman, Sanford J., and Guy Laroque, 1990, Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods, *Econometrica* 58, 25–51.
- Grossman, Sanford J., and Robert J. Shiller, 1982, Consumption correlatedness and risk measurement in economies with non-traded assets and heterogeneous information, *Journal of Financial Economics* 10, 195–210.
- Hansen, Lars P., 1982, Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029–1054.
- , John Heaton, and Erzo Luttmer, 1995, Econometric evaluation of asset pricing models, *Review of Financial Studies* 8, 237–274.
- Hansen, Lars P., and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.
- , 1997, Assessing specification errors in stochastic discount factor models, *Journal of Finance* 52, 557–590.
- He, Hua, and David M. Modest, 1995, Market frictions and consumption-based asset pricing, *Journal of Political Economy* 103, 94–117.
- Heaton, John C., and Deborah J. Lucas, 1995, The importance of investor heterogeneity and financial market imperfections for the behavior of asset prices, *Carnegie-Rochester Conference Series on Public Policy* 42, 1–32.
- Lucas, Deborah J., 1994, Asset pricing with undiversifiable income risk and short sales constraints: Deepening the equity premium puzzle, *Journal of Monetary Economics* 34, 325–341.
- Lucas, Jr., Robert E., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1426–14.
- Luttmer, Erzo. G. J., 1996, Asset pricing in economies with frictions, *Econometrica* 64, 1439–1467.
- Marshall, David, 1994, Asset return volatility with extremely small costs of consumption adjustment, Federal Reserve Bank of Chicago Working Paper No. WP-94-23.
- Marshall, David A., and Nayan G. Parekh, 1998, Can costs of consumption adjustment explain asset pricing puzzles?, *Journal of Finance*, forthcoming.
- Weil, Philippe, 1989, The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* 24, 401–421.

Appendices

A Construction of the Data

The total quarterly real non-durable, durable, services, and total consumption series, the deflators for each of the three components series, and the population series (GPOP) were extracted from CITIBASE for the 1947:1-1997:4 period.

Monthly VW index returns were obtained from CRSP, and were cumulated to obtain quarterly returns. One month T-Bill returns were taken from the CRSP RISKFREE file. One, two, and three year nominally risk-free rates were computed as the returns to one-, two-, and three-year zero-coupon bonds, computed from the Fama-Bliss data in CRSP.

B Construction of the Coherency and Phase Estimates and Confidence Intervals

To construct the coherency and phase estimates plotted in Figures 1 and 2, the two series x and y (log consumption growth and returns) and are first each subdivided into N non-overlapping subsamples of length n (here, $n = 16$ quarters). Each of these subsamples is then detrended and windowed using a Hamming window. The detrended, windowed subsamples are then fast-fourier transformed to generate $J_x(\omega)$ and $J_y(\omega)$, which are equal to, for x ,

$$J_x(\omega) = n^{-1} \sum_{t=0}^{n-1} x_t e^{-i\omega t}.$$

where each of the subsamples is indexed from $t = 0, \dots, n - 1$. Note that J_x and J_y will be complex. The x and y power-spectral densities and the cross-spectral densities are then defined as:

$$\begin{aligned} P_{xx}(\omega) &= \frac{1}{N} \sum_{k=1}^N \frac{n}{2\pi} J_x^k(\omega) J_x^k(\omega)^* \\ P_{yy}(\omega) &= \frac{1}{N} \sum_{k=1}^N \frac{n}{2\pi} J_y^k(\omega) J_y^k(\omega)^* \\ P_{xy}(\omega) &= \frac{1}{N} \sum_{k=1}^N \frac{n}{2\pi} J_x^k(\omega) J_y^k(\omega)^* \end{aligned}$$

where $J_x^k(\omega)$ denotes the Fourier transform of the k 'th subsample, and $*$ denotes the complex conjugate.

Although P_{xx} and P_{yy} are real valued, in general the cross-spectral density will not be. The *coherency* between the two series is defined as:

$$S_{x,y}(\omega) = \frac{|P_{xy}(\omega)|}{[P_{xx}(\omega)P_{yy}(\omega)]^{1/2}}.$$

and the phase is defined as:

$$\phi_{x,y}(\omega) = \arctan \left(\frac{\text{Im}(P_{xy}(\omega))}{\text{Re}(P_{xy}(\omega))} \right).$$

where $\text{Im}(\cdot)$ and $\text{Re}(\cdot)$ denote the imaginary and real components, respectively. With these definitions,

$$P_{xy}(\omega) = (P_{xx}(\omega)P_{yy}(\omega))^{1/2} s_{x,y}(\omega)e^{i\phi_{x,y}(\omega)}.$$

Finally, confidence intervals for the coherency and phase were calculated using the method described in Bloomfield (1976, Section 9.5). The upper and lower bounds of the 95% confidence intervals are therefore:

$$\tanh \left(\text{arctanh}(P_{xy}(\omega)) \pm \frac{1.96g}{\sqrt{2}} \right)$$

where $g^2 = (2/3)(n/T)$, is a constant based on the Hamming window (see p. 224 of Bloomfield (1976).) The 95% confidence intervals for the phase are:

$$\phi_{x,y}(\omega) \pm \sqrt{\frac{1}{P_{xy}(\omega)^2} - 1}.$$