Habit Formation: A Resolution of the Equity Premium Puzzle?

Christopher Otrok
University of Virginia

B. Ravikumar
Pennsylvania State University

Charles H. Whiteman*
University of Iowa

First Version: January 1998

This Version: March 2001

Abstract

We explore how the introduction of habit preferences into the simple intertemporal consumption-based capital asset pricing model "solves" the equity premium and risk-free rate puzzles. While agents with time-separable preferences care only about the overall volatility of consumption, we show that agents with habit preferences care not only about overall volatility, but also about the temporal distribution of that volatility. Specifically, habit agents are *much* more averse to high-frequency fluctuations than to low-frequency fluctuations. In fact, the size of the equity premium in the habit model is determined by a relatively insignificant amount of high-frequency volatility in U.S. consumption. Further, the model's premium and returns are very sensitive to changes in characteristics of the stochastic process for consumption, changes that have been dramatic during the 20th century. The model also carries counterfactual implications for the equally dramatic changes in the equity premium and risk-free rate observed over the last 100 years.

---

* Corresponding author: Charles H. Whiteman, Department of Economics, University of Iowa, Pappajohn Business Building, Iowa City, IA 52242. We thank George Constantinides, Dave DeJong, Lars Hansen, Beth Ingram, Rody Manuelli, Mark Watson and an anonymous referee for helpful comments. We gratefully acknowledge the support provided for this research by the National Science Foundation under grants SES-0082237 and SES-0882230. Email: cmo3h@virginia.edu; ravikumar@psu.edu; whiteman@uiowa.edu.

I. Introduction

Is consumption habit-forming? If so, it might be possible to account simultaneously for the high equity premium and low risk-free rate of return observed in historical data for the U.S. Mehra and Prescott (1985) argued that the observed premium of equity returns over bond returns was puzzling in the context of an economy populated by individuals with time-separable constant relative risk aversion preferences. Constantinides (1990), in contrast, showed that if consumption levels in adjacent periods are complementary in agents' preferences, these puzzles could be resolved. Our aim in this paper is to explore the nature of this resolution. To do so, we employ *spectral* utility functions that help us examine an agent's attitude toward the temporal distribution of consumption volatility. We then show that if the representative agent's preferences exhibit intertemporal non-separabilities (such as habit), the asset-pricing model has important implications for the equity premium and the risk-free rate of return.

Spectral utility functions decompose agents' preferences for consumption smoothness into preferences for smoothness by frequency. They are calculated by first decomposing consumption time series into orthogonal components that have all volatility concentrated in given frequency bands. Then expected utility is calculated for each orthogonal component. For time-separable preferences, spectral utility is flat because in such cases agents care only about overall volatility (and, of course, higher moments) in consumption and not about the temporal distribution of volatility. But for habit preferences the spectral utility function is not flat: agents are less averse to low-frequency (persistent) fluctuations than to high-frequency (less persistent) fluctuations. Moreover, as overall volatility increases, habit agents disproportionately eschew high-frequency fluctuations.

The spectral utility calculations suggest that the risk premia demanded by habit agents to hold risky securities will depend on whether the securities' payoffs have high- or low-frequency volatility. Indeed, we find that without changing overall volatility, making consumption volatility less persistent can lead to dramatic increases in the equity premium in the habit model. For instance, such a shift, designed to be modest relative to historically observed changes in consumption, can lead to increases in the equity premium in excess of 1600 basis points.

We also show that the size of the equity premium in the habit model is determined by a relatively insignificant amount of volatility over short time horizons. For instance, if high frequency volatility is held fixed, even an immodest *tripling* of overall consumption volatility may be accompanied by a *decline* in the equity premium so long as consumption is made more persistent.

That the equity premium in the habit model is driven by high-frequency volatility poses a problem for the model, since the distinguishing characteristic of U.S. consumption is that it is exceedingly *smooth*, i.e., the consumption volatility is concentrated at low frequencies. However, as we show, the habit model is extremely sensitive to changes in even the tiny amount of high-frequency consumption volatility observed in U.S. data. Moreover, this sensitivity of the habit model suggests that the model's predictions might not be robust to historically observed changes in the properties of consumption growth. To investigate this, we estimate the observed changes in consumption growth moments using 40-year rolling samples, and find that volatility has declined by a factor of 4, while the first-order autocorrelation has increased from about –0.5 to 0.4. Using preference parameters calibrated to reproduce the average equity premium and risk-free rate over the last 100 years together with the consumption growth moments from the 40-year rolling samples, we show that the habit model predicts a dramatic 2500-basis-point decline in the equity premium and a dramatic 1600-basis-point increase in the risk-free rate. In stark contrast, the observed equity premium has risen by about 200 basis points, and the observed risk-free rate has declined by a similar amount.

We produce the above results with a simple yet popular version of habit preferences where the habit stock is a linear function of just the previous period's consumption. This version of habit is illustrative and makes the implications of the model transparent. Our main results, however, are robust to alternative specifications of habit. We replicate our results for two other specifications: (i) a multiple lag internal habit formation similar to Heaton (1995) and an external habit formation similar to Campbell and Cochrane (1999).

II. Models of the Equity Premium and Risk Free Return

We utilize the Lucas (1978) "tree" model as our foundation for general equilibrium asset pricing. In this economy, there is a single tree that yields an exogenous stochastic flow of fruits, denoted by d_t at time t. The representative agent in this economy has preferences described by

$$E_o \sum_{t=0}^{\infty} \beta^t u(c_t)$$

where E_0 denotes conditional expectation given information at time 0, $u(\cdot)$ denoted the period utility, c_t denotes consumption at time t and $\beta \in (0,1)$ is the discount factor. There is a competitive market for trading claims to the tree's fruits. The measure of agents and the measure of outstanding claims are each normalized to one, with the representative agent holding the single claim to the tree's fruits. With p_t denoting the price of one claim and s_t denoting the agent's shareholdings at time t, the agent's budget constraint is given by:

$$c_t + p_t s_{t+1} \leq (p_t + d_t) s_t .$$

The agent's first-order conditions for choosing the optimal consumption and shareholding sequences are:

(1) $$u'(c_t) p_t = \beta E_t u'(c_{t+1})(p_{t+1} + d_{t+1}), \qquad t \geq 0.$$

In equilibrium all of the fruits are consumed each period and there is no other source of the consumption good, so $c_t = d_t$ for all t. The equilibrium prices are then determined as stationary functions of the state: $p_t = p(c_t)$. This method can be used to price assets with different payoff structures as well. For instance, the time-t price, p_t^f, of a one-period bond that pays one unit of consumption in period t+1 must satisfy

(2) $$u'(c_t)p_t^f = \beta E_t u'(c_{t+1}), \qquad t \geq 0.$$

Thus, the return on the risky asset is given by:

(3) $$R_{t+1} = \frac{p_{t+1} + c_{t+1}}{p_t}$$

and the return on the riskless asset is given by:

(4) $$R_{t+1}^f = \frac{1}{p_t^f}.$$

The 'equity premium' is the difference between the two returns.

Mehra and Prescott (1985) studied whether this model is consistent with the observed equity premium in the U.S. They assumed a specific functional form for u(·) and a specific stochastic process for consumption. Period utility was of the constant relative risk aversion class, so preferences were given by:

$$E_o \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma}, \qquad \sigma > 0.$$

(The $\sigma = 1$ case will be interpreted as logarithmic.) Consumption growth was approximated by a two-state Markov process with a symmetric transition matrix. The three parameters of this process (the values of consumption growth in the 'good' state and the 'bad' state, and the probability of changing from one state to another) were chosen to match the mean, standard deviation and first-order autocorrelation of annual growth of real per-capita consumption of nondurables and services from 1889-1978. The return on the risky asset was measured by the dividend-inclusive *ex-post* real return on the S&P500 and the return on the riskless asset was the *ex-post* real risk-free return on 3-month Treasury bills (relative to the consumption deflator for nondurables and services—see Shiller, 1989).

Given values for β and σ, Mehra and Prescott showed that the price of the risky asset is proportional to consumption, with the constant of proportionality depending on the state. They also derived closed-form solutions for the returns on the risky and the risk free assets. For $\beta \in (0,1)$ and $\sigma \in (0,10]$, they found that the average equity premium implied by the model is too small relative to the data.

The equity premium 'puzzle' can be solved by using very large values of risk aversion (e.g., $\sigma = 50$), but this was regarded as implausible. Moreover, as has been argued by Weil (1989), even with high degrees of risk aversion, it is not possible with time-separable preferences to account simultaneously for the large premium of stock returns over bond returns and the low return on essentially risk-free short-term Treasury securities. Indeed, when we use consumption and return data from 1889-1992, $\beta = 0.96$ and $\sigma = 2$, the model implies an average equity premium of 0.25%, and a risk free rate of 7.44%, whereas their counterparts in the data are 6.87% and 0.90%.

There have been many attempts in the literature to explain the risk-free rate and equity premium puzzles, but one of the most popular was advanced by Constantinides (1990). He argued that consumers are habitual, in that consumption values in adjacent periods are complementary.[1] That is, the preferences of consumers in a discrete-time version of his model are given by:

(5)
$$E_o \sum_{t=0}^{\infty} \beta^t \frac{[(1+\delta(L))c_t]^{1-\sigma}}{1-\sigma},$$

where $\delta(L)$ is a polynomial in the lag operator L. The interpretation of these preferences depends on the sign of the lag coefficients in the polynomial $\delta(L)$. When the lag coefficients are all negative, the preferences exhibit habit-persistence. If the coefficients are all positive, then the utility function is said to display durability. When the coefficients in $\delta(L)$ are zero, the preferences are time-separable.

With habit-persistence preferences, current utility depends on the level of consumption in the current period relative to the level of consumption in previous periods. The central idea is that once the agent consumes at a certain level, he or she gets accustomed to that level of consumption. With durability, the agent has a technology that permits the transformation of goods purchased in one period into a flow of consumption services in future periods.

Using a relatively low 'risk aversion' parameter coupled with a particular habit function, Constantinides' calculations suggested general consistency among the mean and standard deviation of consumption growth, mean risk-free rate, and mean equity premium observed over 1889-1978. In our version of these calculations, we initially utilize a simple one-lag habit polynomial ($\delta(L) = \delta L$). We calculate asset prices using the simulation method described in Judd (1998). The method requires specifying a parametric time series model of consumption, and then simulating the consumption process repeatedly to evaluate the expectation in equations (1) and (2). We estimate an AR(1) time series process for consumption growth using data from 1889-1992. Additional details on the price calculation can be

[1] Other explanations of the equity premium puzzle using time-non-separable preferences include: Abel (1990), Heaton (1995), Boldrin, Christiano and Fisher (1997), Campbell and Cochrane (1999), and Jermann (1998). Bansal

found in the Appendix. Given one-lag habit and $\beta = 0.955$, we set $\delta = -0.615$ and $\sigma = 0.8$; together with the consumption growth process, these values of δ and σ uniquely enable us to reproduce the average equity premium and risk free rate over the period 1889-1992.[2]

III. A Spectral View of Habit and Time-Separable Preferences

How does the habit model deliver the observed equity premium and the risk fee rate with low σ?[3] There is one key feature of the habit model that accounts for this: agents with habit preferences care not only about the overall volatility of consumption but also about the temporal distribution of that volatility.

The next question, then, is how does habit reflect preferences for the temporal nature of consumption fluctuations? A useful way to measure the temporal characteristics of consumption itself is the spectrum, which provides a frequency-by-frequency decomposition of the *variance* of a time series. To measure how agents feel about consumption fluctuations we use spectral utility functions, which measure *expected utility* frequency by frequency.

III.1 Spectral Utility Functions

Spectral Utility is the level of expected utility associated with a particular frequency of consumption volatility; it is a function that maps frequencies into real numbers. With time-separability and quadratic utility, expected utility can be represented in a two-dimensional plot of the mean and variance of the consumption time series. But with time-non-separability, the frequency of variance matters and a third dimension is required. We refer to that third dimension as spectral utility. What the spectral utility function does is assign a number to the level of expected utility at each frequency of consumption volatility. In the special quadratic case, spectral utility can be calculated analytically (see Whiteman, 1985, 1986, for derivations, and Taub, 1989, for applications). Calculating spectral utility for more general cases, such as constant relative risk aversion or habit preferences, requires numerical approximation. We calculate spectral utility functions for preferences of the form in (5) by decomposing a time series for consumption into different frequency components, and then computing the expected utility of each frequency component.

The decomposition of consumption is accomplished by applying band-pass filters to the time series. When a time series is put through such a filter, the filtered series will have all its fluctuations

---

and Coleman (1996) propose a monetary explanation that is based on time-separable preferences. Kocherlakota (1996) provides a survey of the literature.

[2] All of the computations in the paper were also done using the two-state Markov process of Mehra and Prescott. The results were qualitatively and quantitatively similar to those reported here.

[3] Unlike the case of time-separable preferences, it is misleading to think of σ as measuring risk aversion in the habit model since the utility over consumption gambles at any point in time depends on past consumption as well as current wealth.

concentrated in one frequency band. In other words, the filtered time series will have a power spectrum (nearly) equal to zero outside of the desired frequency band.[4] Different band-pass filters isolate different components of the time series.

To calculate spectral utility, we begin by calculating band-pass filters for a complete set of narrow frequency bands. For example, we split the interval $[0,\pi]$ into 32 equal bands, and let the associated band-pass filters $b_j(L)$ be such that $b_j(\omega) = 1$ for $\omega \in [j\pi/32,(j+1)\pi/32]$ for $j = 0,1,\ldots 31$ and zero otherwise. After calculating the filters, the procedure for calculating spectral utility is a three-step process: first, simulate a time series realization of consumption; second, apply the 32 band-pass filters to the realization to isolate the consumption fluctuations by frequency band; third, calculate realized utility in each frequency band as the utility of the consumption time series associated with that frequency band. Expected utility in each frequency band is the mean of realized utility in the frequency band across a large number of replications of the three steps.

For example, consider an iid time series realization of consumption. A band-pass filter for the 16-period cycle gives us a particular component of the iid time series. We evaluate the utility associated with this component. We can evaluate similar utilities associated with other components derived from band-pass filters for other frequency bands. This gives us realized utilities, frequency by frequency, for a particular drawing of consumption. To calculate expected utility frequency by frequency, we need many realized utilities in each frequency band. We achieve this by repeating the above procedure for many simulated drawings from the consumption process.[5] If we denote simulation n (of N) of the consumption time series of length T by $\{c_{n,t}\}_{t=1}^{T}$, realized utility for frequency band j is approximated by

(6) $$U_n(j) = \sum_{t=1}^{T} \beta^t u(b_j(L)c_{n,t}) .$$

Expected utility for frequency band j is thus given by

(7) $$EU(j) = \frac{1}{N}\sum_{n=1}^{N} U_n(j) .$$

Note that the band-pass filtered consumption series constitute an additive decomposition of overall consumption. Spectral utility does not share this property unless agents are risk neutral; that is, spectral utility is not an additive decomposition of overall expected utility. Of course, the decomposition need not be additive to help us understand the properties of temporally dependent preferences. Furthermore, spectral utility is not the same as the spectrum of realized utility. The former is the expected utility associated with consumption fluctuations at specific frequencies, whereas the latter is an orthogonal decomposition of the *second* moment of the utility process. While the mean of realized utility, expected

---

[4] See Baxter and King (1999) and Christiano and Fitzgerald (1999) for more details on band-pass filtering.

utility, is the object of usual economic interest, the second moment of the utility process has no apparent economic interpretation.[6]

Figure 1 displays spectral utility for serially uncorrelated *lognormal* consumption for two cases: the time-separable case and a "habit" case with δ = -0.615. In both cases, β = 0.955, σ = 0.8. (For the "durability" case, e.g., δ = 0.5, spectral utility slopes upward.) An agent with time-separable preferences is indifferent to various temporal distributions of volatility as long as the overall volatility remains the same. In this case, spectral utility is flat: two different consumption processes, one with all its volatility concentrated at low frequencies and another with the same volatility concentrated at high frequencies, yield the same expected utility. The figure also illustrates an important feature of the spectral utility associated with habit preferences. As Orwell might have put it, some variances are more equal than others: for habit preferences, spectral utility slopes downward (see Figure 1b). The two consumption processes yield different expected utilities for an agent with habit preferences: the agent would prefer the consumption process with volatility concentrated at low frequencies to the less persistent, high-frequency volatility consumption process.[7]

Figure 1: Spectral Utility Functions

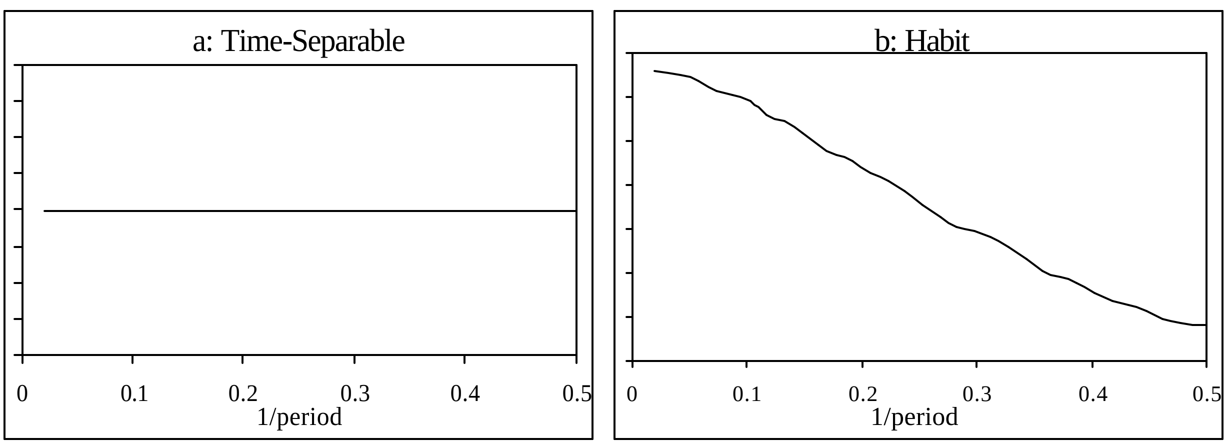

A second feature is that some variances are *disproportionately* more equal than others. This is illustrated in panels a and b of Figure 2. These panels display expected utility by frequency for several different values of overall consumption variance, where volatility at each frequency is increased by the

---

[5] Additional details of these computations can be found in Otrok (2001).

[6] The utility process is the sequence $\sum_{k=0}^{\infty}\beta^k u(c_{t+k})$ for t = 1, 2, … ∞. The mean of this process is expected utility. The spectrum associated with the sequence is a decomposition of the variance of the utility process.

[7] For a serially correlated consumption process, the variance will not be distributed equally across all frequencies, so the interaction between the consumption spectrum and preference nonseparabilities will produce more complicated spectral shapes than those in Figure 1.

same percentage. Panels c and d show expected utility as a function of consumption variance for two frequencies, a "long-run" frequency corresponding to 50-period cycles, and a "short-run" frequency corresponding to 2-period cycles. For time-separable preferences, the drop in expected utility is the same for both frequencies whereas for habit preferences, high-frequency spectral utility drops more rapidly than low-frequency spectral utility (panel d).

Figure 2: Spectral Utility Functions for Different Consumption Volatilities

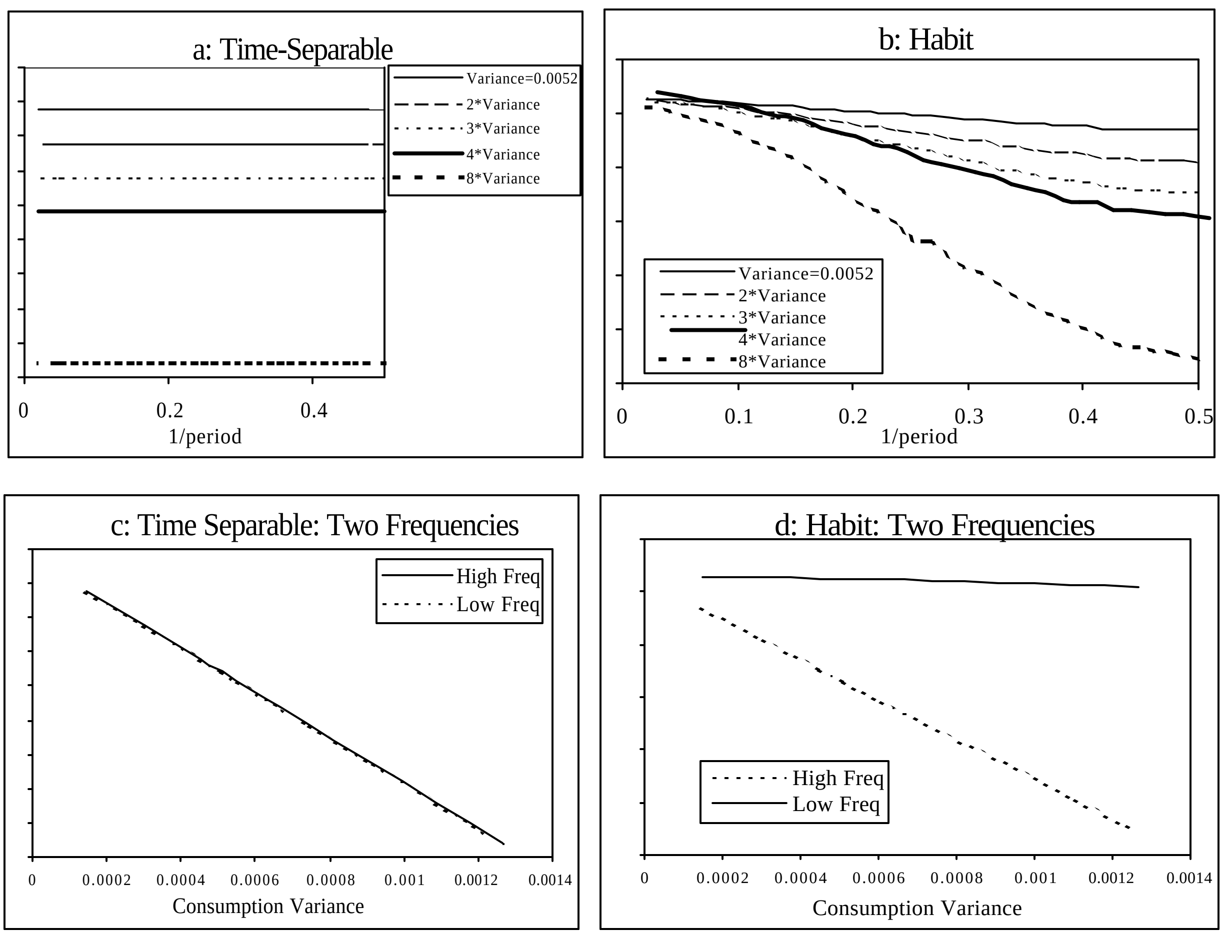

As is evident from Figure 2b, this habit agent is indifferent to an eight-fold increase in volatility when the volatility is shifted from a high frequency (1/period = 0.5) to a low frequency (1/period = 0.05). That is, faced with increasing variance, the habit consumer wants it concentrated more and more at lower frequencies; when it is not, he or she would demand a higher premium to bear the associated risk.

It is clear from Figures 1 and 2 that consumption processes with different spectral characteristics will have differing effects on the compensating risk-premia demanded by individuals with habit preferences. The quantitative implications of this feature are illustrated in the next subsection.

III.2 Implications for the Equity Premium

The calculation of the equity premium and risk-free rate in Section II requires a consumption growth process. We first consider three processes for the *level* of consumption that have the same overall volatility but different temporal distributions of volatility. In Figure 3, we have plotted three spectra for log consumption where the autocorrelation of consumption *growth* is 0.3, 0, and -0.3. (Throughout, we hold the mean consumption growth rates fixed across the three specifications at the 1889-1992 sample average of 1.7 percent. This is accomplished by adjusting the constant term in the autoregression.)

Figure 3: Log Consumption Spectra: Constant Level Variance

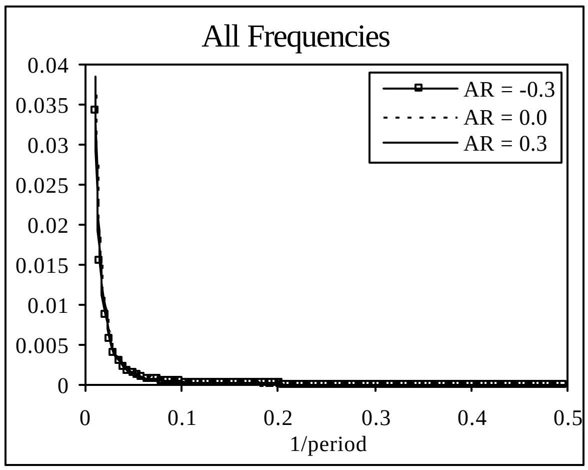

These spectra are derived by starting with an AR(1) process for consumption growth and then applying the $(1-L)^{-1}$ filter to the consumption growth spectrum; i.e., dividing the growth spectrum by $1-e^{-i\omega}$ except at $\omega = 0$. The *level* variance is held constant by adjusting the innovation variances in the AR(1) growth processes. In all three spectra in Figure 3, the standard deviation of the log level of consumption is 0.0690.[8] The resulting equity premia and the risk-free rates for habit preferences with $\delta = -0.615$, $\beta = 0.955$ and $\sigma = 0.8$ are given in Table 1.

[8] This standard deviation of the log level of consumption arises from an iid consumption growth process (as in Constantinides, 1990) with standard deviation 0.03, a historically reasonable value. The standard deviation of the log level was calculated as the square root of twice the spectrum integrated over the range $[0.02\pi,\pi]$ (a range corresponding to periodicities of up to 100 years) involving 1024 ordinates.

Table 1: Habit Model Asset Returns with Constant Consumption Variance

Cons. Growth Auto. Corr.	0.3	0.0	-0.3
Equity Premium (%)	0.67	3.76	16.8
Risk Free Rate (%)	5.71	3.31	-6.24
Cons. Growth Std. Dev.	0.0227	0.03	0.0397

It is clear that the equity premium is very sensitive to changes in the temporal distribution of consumption volatility—when the autocorrelation of consumption growth changes from 0.3 to -0.3, the equity premium increases by more than 1600 basis points. This is despite the fact that the three level spectra in Figure 3 do not appear to be very different (though of course the growth rate spectra are).

To examine further the role of the temporal distribution of volatility, consider three consumption processes where the *high-frequency volatility* is the same. Specifically, we construct three consumption level processes as explained earlier, but instead of holding the overall level variance constant, we hold constant the variance in the 2 to 3 year range (1/period = 0.5 to 0.33). Despite a reduction of 60% in the log consumption standard deviation (0.1209 to 0.0411) in moving from growth autocorrelation of 0.3 to -0.3, the equity premium actually rises because high-frequency volatility has become relatively more important (see Table 2). This suggests that the differences between the three processes at low frequencies have very little impact on the equity premium for the habit model—the sensitivity of equity premium in Table 1 is the result of (visually indistinguishable) differences in the high frequency part of the spectra in Figure 3.

Table 2: Habit Model Asset Returns with Constant High Frequency Consumption Variance

Cons. Growth Auto. Corr.	0.3	0.0	-0.3
Equity Premium (%)	3.19	3.76	3.90
Risk Free Rate (%)	3.72	3.31	3.20
Overall Cons. Level Std. Dev.	0.1209	0.0690	0.0411
Cons. Growth Std. Dev.	0.0372	0.03	0.0263

To complete the description of how volatility at different frequencies affects the equity premium, we consider three consumption processes where the *low-frequency volatility* is the same. The resulting equity premia and risk-free rates are now given in Table 3.

Table 3: Habit Model Asset Returns with Constant Low Frequency Consumption Variance

Cons. Growth Auto. Corr.	0.3	0.0	-0.3
Equity Premium (%)	0.71	3.76	18.89
Risk Free Rate (%)	5.66	3.31	-7.91
Overall Cons. Level Std. Dev.	0.0686	0.0690	0.0699
Cons. Growth Std. Dev.	0.0225	0.03	0.0405

Even though the overall level volatility is almost the same, the equity premium increases by more than 1800 basis points when the consumption growth autocorrelation falls from 0.3 to -0.3.

These calculations illustrate that the *temporal* nature of consumption volatility is at least as important as the *magnitude* of consumption volatility. Huge increases in the equity premium can be obtained (as in Table 1) by relocating volatility from low to high frequencies. Moreover, once high-frequency volatility is held fixed (Table 2), a tripling of overall variance does not increase the equity premium, while tiny changes (Table 3) in overall variance concentrated at the high frequencies produce enormous changes in the equity premium. This is parametric sensitivity on a grand scale.

Figure 4 illustrates just how grand the scale really is. The figure depicts two deterministic time series associated with the “AR = -0.3” and “AR = 0.3” cases, constructed from a constant-low-frequency-consumption-variance experiment like that underlying Table 3. These time series are weighted sums of sine waves at various frequencies, with weights given by the square roots of the heights of the spectra. An agent would demand a greater premium to compensate for the fluctuations about mean consumption represented by the (slightly) more volatile “AR = -0.3” consumption process than to compensate for the “AR = 0.3” process despite the fact that the two consumption processes (the two “High+Low Frequency Components”) are almost the same. As indicated in Table 3, the equity premium demanded by the habit agent to hold the “AR = -0.3” consumption process is more than 1800 basis points higher than the one demanded to hold the “AR = 0.3” consumption process, and the reason is the difference between the high frequency components that are each practically invisible.

Figure 4: Components of Consumption Time Series

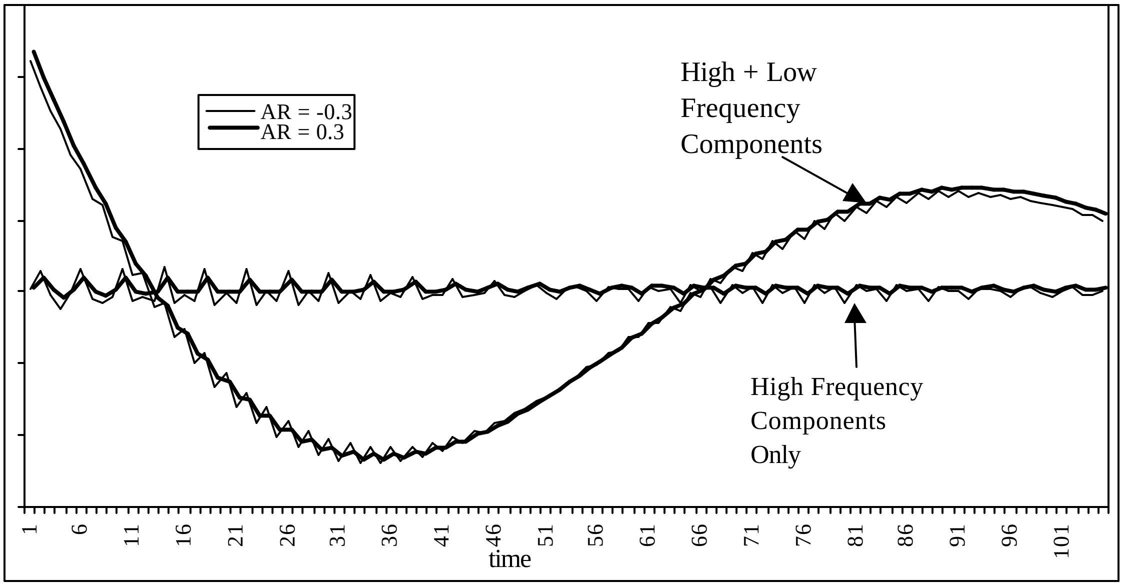

In the next section, we use the two models (habit and time-separable) together with *observed* changes in the consumption process to calculate the equity premium and risk-free rate. We then examine whether the calculated values are consistent with their empirical counterparts.

IV. Confronting the Changing Characteristics of the Equity Premium and Risk-Free Rate

IV.1 Changes in Consumption Growth in the U.S.

The changing nature of the consumption growth process over time is reflected in the consumption growth moments in the pre-war (1889-1929) and post-war depicted in Figure 5. To emphasize how the *estimates of the moments* evolve through time, the figures are calculated using 40-year rolling samples. In each case, the date on the horizontal axis corresponds to the ending period of the sample. Each of the three panels in the Figure reflects remarkable secular movements in consumption growth moments over the past 100 years. (See also Golob (1992) for evidence on the changes in the consumption process.)

Figure 5: Evolution of Consumption Growth Moments

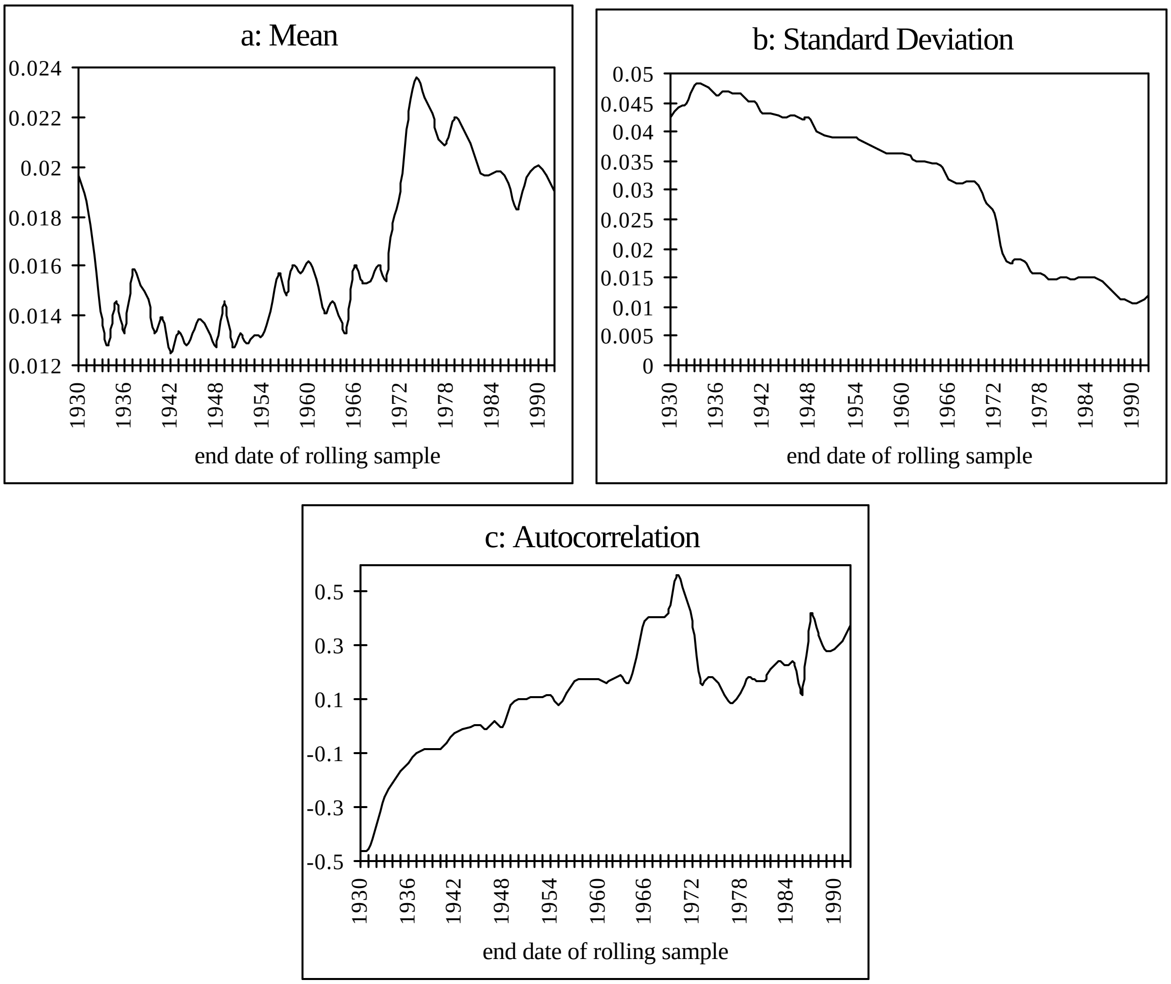

Figure 5a, for example, reflects the enormous effect the Great Depression had on average consumption growth in the U.S. In the 40-year sample ending just prior to the Depression's onset, mean consumption growth was nearly 2% per year; in all of the ensuing 40-year samples that include it, mean consumption growth is nearly a third smaller. Moreover, in the later samples that do not include the Depression, mean consumption growth is much higher.

Figures 5b and c indicate that consumption growth has become more predictable during the 20th century. Figure 5b, for example, shows that the standard deviation of consumption growth has declined over time, indicating that consumption growth has become more predictable in the sense that it has become more concentrated around its mean. Figure 5c documents that autocorrelation in consumption growth has risen during the 20th century. This also means that consumption growth has become more predictable—consumption growth one year ahead has become increasingly similar to current consumption growth.[9]

IV.2 Evolution of the Equity Premium and the Risk-Free Rate: Data vs. Theory

The implications of the changing consumption growth moments for the time-separable model with $\beta = 0.96$ and $\sigma = 2.0$ are illustrated in Figure 6. Panels a and b of Figure 6 display the mean equity premium and the mean risk-free rate produced by the time-separable model (thin curve, "TS") using the corresponding consumption growth moments estimated for the appropriate sub-sample; the panels also display the observed mean equity premium and risk-free rate (thick curve). The calculations proceed as if at each date, agents treat the current "rolling" estimates of the consumption growth moments as true values, though the calculations do not admit the possibility that agents take into account the apparently changing nature of the consumption process. The figures are best viewed as indicating what the *average* equity premium and risk free-rates were historically (using rolling 40-year samples) compared to what the model generates as the mean equity premium and risk-free rate using the then-prevailing estimates of the parameters of the consumption growth process.

For example, consider the time-separable model's prediction and the data in Figure 6 for the end date 1930. The mean, standard deviation and autocorrelation of consumption growth for the 40-year subsample from 1890 to 1930 are 1.96 percent, 4.26 percent and -0.46 respectively (as illustrated in Figure 5 for the end date 1930). The representative agent in the time-separable model uses these consumption growth moments to decide on consumption and asset-holdings. These decisions, in equilibrium, imply a mean equity premium of 0.19 percent and a mean risk-free rate of 6.05 percent. For the period 1890-1930, the mean equity premium is 5.67 percent and the mean risk-free rate is 2.79 percent. These numbers — the model's implications and the data on mean equity premium and risk-free rate — are illustrated in Figure 6 for the end date 1930.

It is clear from Figure 6 that the considerable changes in the characteristics of consumption growth do not produce noticeable changes in the mean equity premium and risk-free rate produced by the time-separable model. In fact, there is much more "action" in the equity premium and risk-free rate in the

[9] The interpretation of these facts adopted in the next subsection is that they do represent changes in the evolution of U.S. consumption. An alternative interpretation is that the data are riddled with measurement error that has declined in severity over time. For example, the standard deviation and autocorrelation estimates in Figure 5 are consistent with actual consumption being a random walk, and measured consumption being actual consumption plus a white noise whose variance has declined over time. If agents observe actual consumption, with unchanging stochastic properties, the model does not predict the dramatic changes in the mean equity premium and risk-free rate of the next section. Of course, such a separation of the information sets of agents and researchers studying them would require reinterpretation of other empirical results involving the equity premium, and indeed the puzzles themselves.

data than in the model. More importantly, in all subsamples, the model produces a mean equity premium that is far too low, and a risk-free rate that is far too high.

Figure 6: Equity Premium and Risk Free Rate Over Time: Data vs. the Time-Separable Model

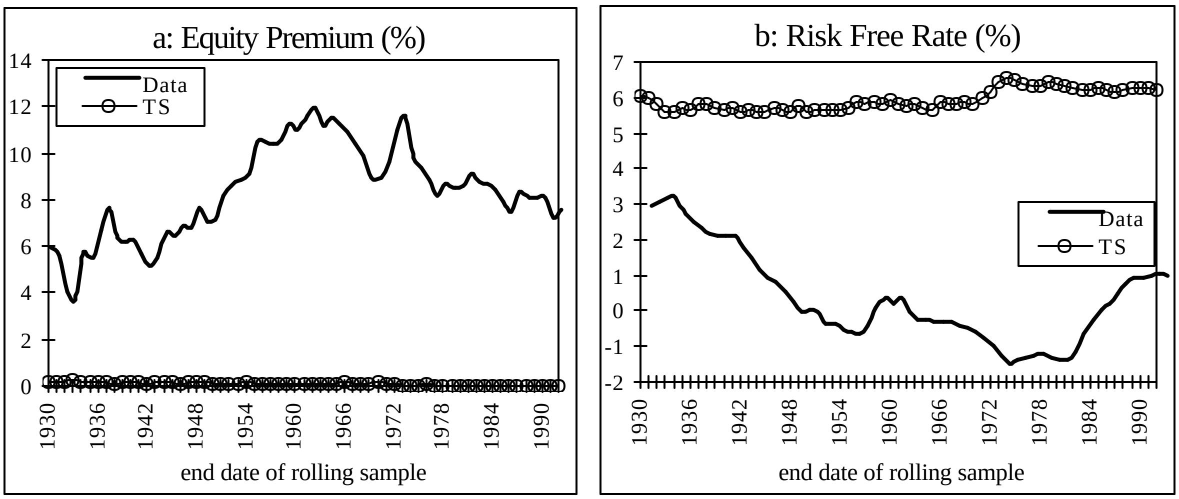

Panels a and b of Figure 7 illustrate similar calculations for the habit model with δ = -0.615, β = 0.955 and σ = 0.8. (Recall from Section II that these parameter values permit the habit model to match the mean equity premium and risk-free rate for the 1889-1992 sample.) The habit model is quite sensitive to changes in the characteristics of consumption. For example, using the parameters that enable a perfect match for the full sample, Figure 7a shows that the model overstates the equity premium by more than 2000 basis points (for a sample ending during the Great Depression), and understates it by nearly 1000 basis points (for a 40-year sample ending in the early 1970's.). Figure 7b shows that the model is about 1400 basis points too low for the risk-free rate in a sample ending in the early 1930's, and 800 basis points too high for a sample ending in the early 1970's. The figure makes it clear that the temporal changes in the mean equity premium and the mean risk-free rate predicted by the habit model, in contrast to the time-separable case, are much more dramatic than in the data.

Figure 7: Equity Premium and Risk Free Rate Over Time: Data vs. the Habit Model

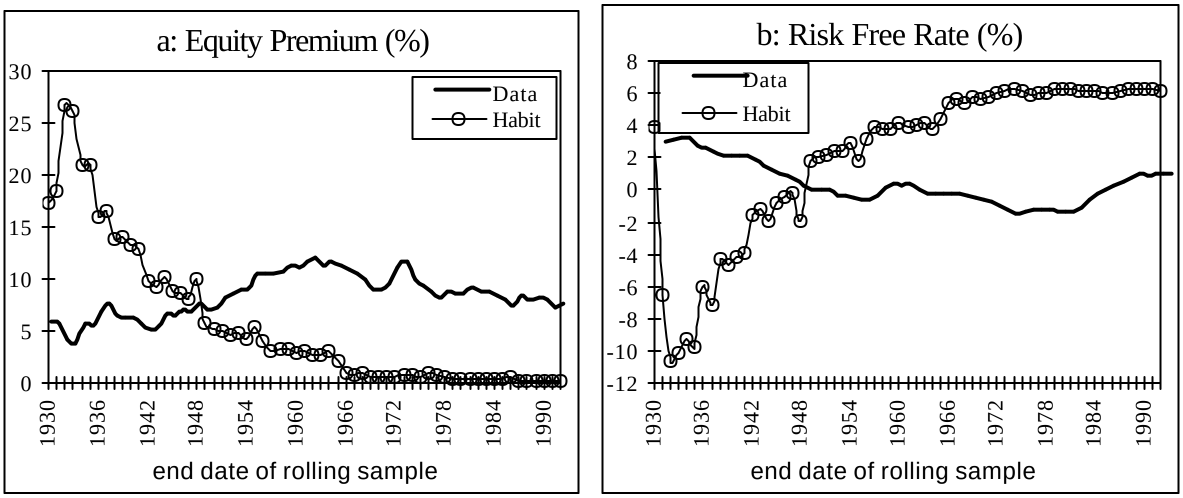

The dramatically counterfactual predictions of the habit model are robust along several dimensions. First, the nature of the results in this section does not change if, instead of the rolling samples, we use expanding samples with a fixed start date or ever shrinking samples with a fixed end point.[10] Second, when we relax the assumption that consumption growth equals dividend growth and allow consumption growth and dividend growth to be parameterized separately, the results are nearly identical to those obtained in this section. Finally, as we will see in section V, using more general forms of habit formation does not significantly change the results in Figure 7.

IV.3 Mechanics of the Time-separable and Habit models

The behavior of the equity premium and risk-free rate in the time-separable and habit models in Figures 6 and 7 can be understood with the help of the consumption growth moments in Figure 5 and the spectral utility diagrams in Figures 1 and 2. Figures 8 and 9 record the effects of changes in the standard deviation and autocorrelation of consumption growth. Risk aversion implies that as the standard deviation of consumption growth falls over time, the premium demanded to hold risky equity will fall and the risk-free rate will rise. This occurs in both models (Figures 8a-b and 9a-b), though the effect in the time-separable model is miniscule compared to that in the habit model.

---

[10] With expanding samples and a fixed start date, there is a kind of myopic learning going on—in effect, the representative agent is running a time series regression on the dividend (consumption) process, and at each date, uses the most recent estimates of the process parameters to calculate the expected present value of future dividends that determines current consumption behavior. As indicated above, this learning is myopic in the sense that the agent does not consider the possibility that the consumption process is changing, and merely believes himself or herself to be getting better estimates of fixed but unknown parameters.

Figure 8: Consumption Growth Moments and Asset Returns: Time-Separable Model

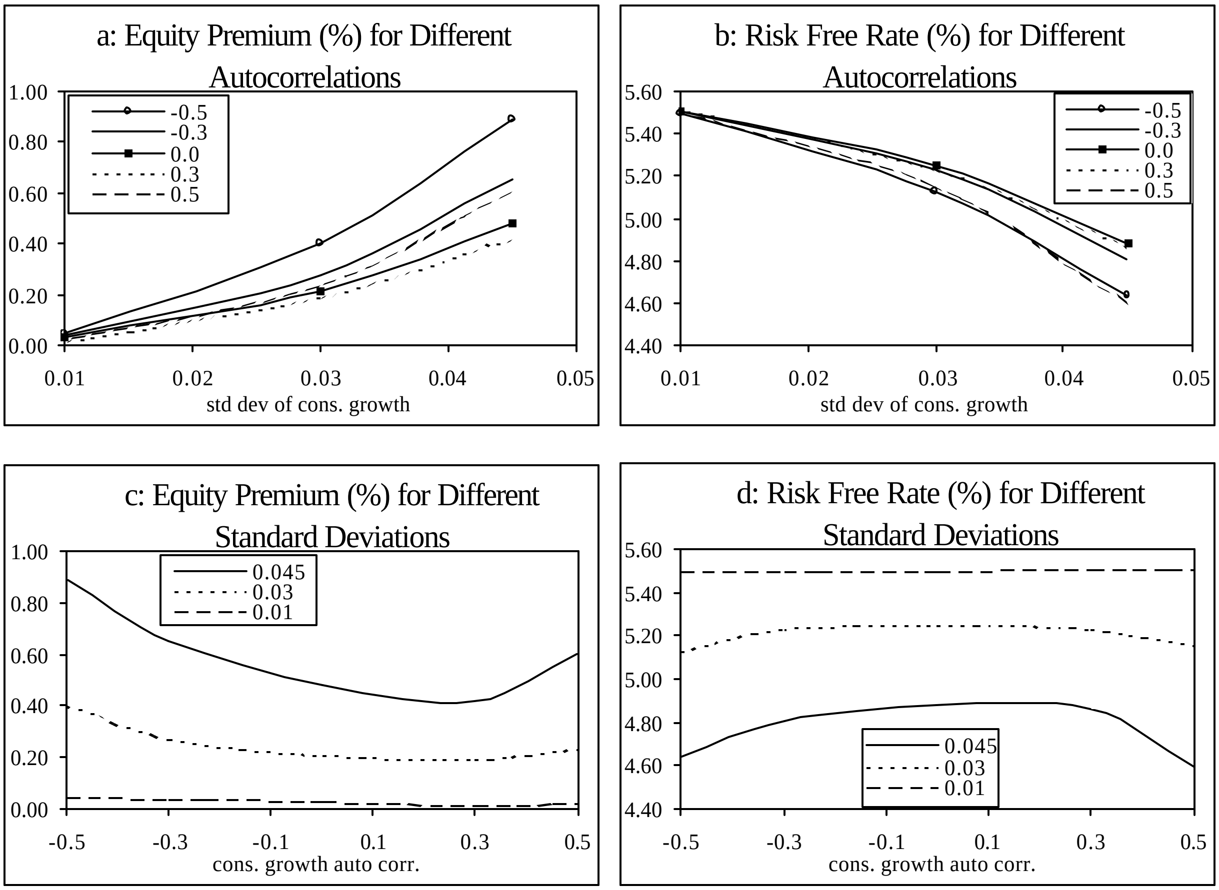

The two models behave quite differently in response to changes in consumption growth autocorrelation. In the time-separable model, any predictability (positive or negative autocorrelation) raises the equity premium and decreases the risk-free rate—though again the magnitudes of the effects are very small (Figures 8c and 8d). In the habit model, in contrast, the rising autocorrelation of consumption growth implies that the consumption spectrum is becoming more peaked and concentrated at lower frequencies. Together with the aversion that habit agents display toward *high* frequency fluctuations, this migration of fluctuations to the lower frequencies causes the equity premium to fall.

The equity premia produced by the habit model (Figure 7a) thus reflect complementary consequences of the reduced standard deviation and increased autocorrelation of consumption growth observed in the sample. Each of these effects works to reduce the equity premium in the model. But from Figure 7a, the unfortunate fact is that the observed equity premium *rises*.

Figure 9: Consumption Growth Moments and Asset Returns: Habit Model

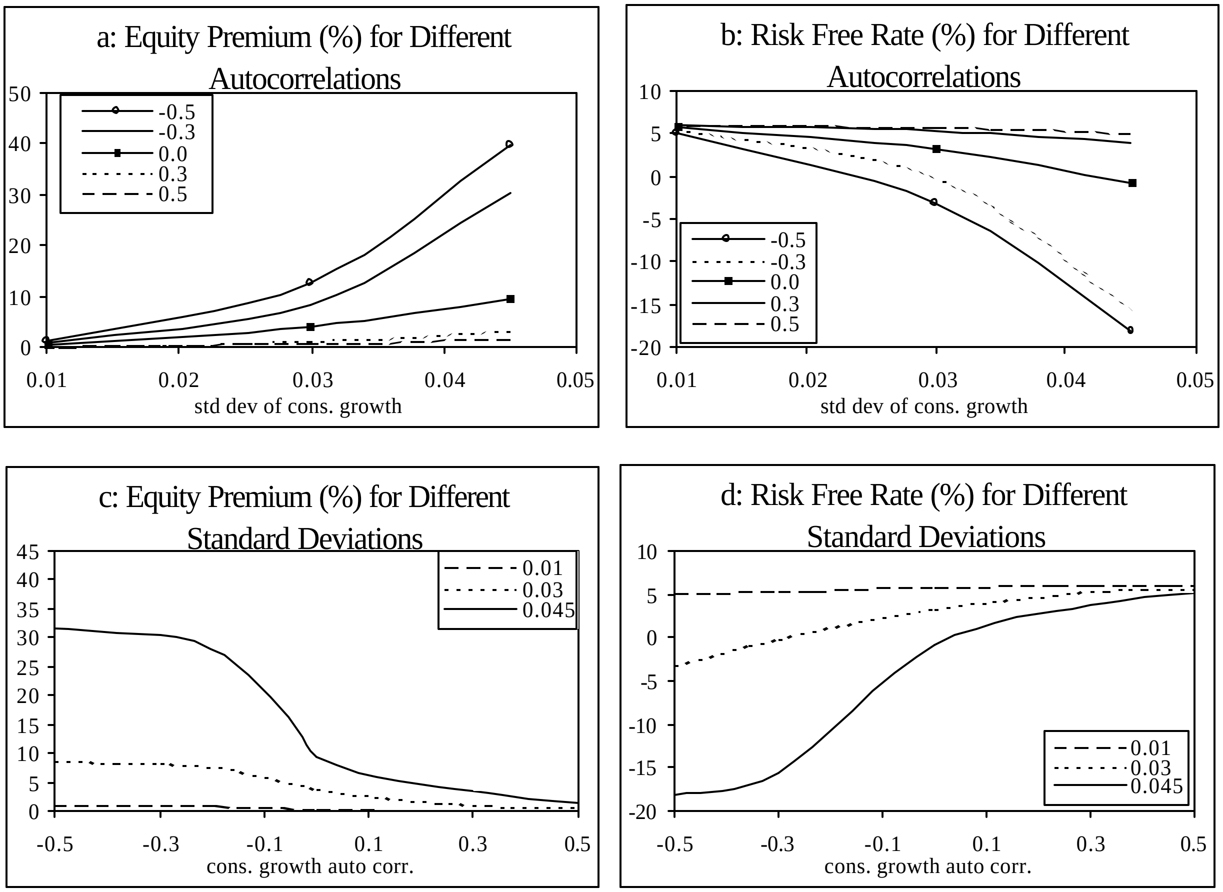

Understanding how the habit model's risk-free rate changes over the sample requires knowledge not just of the changing second moments of consumption growth, but of the mean as well. From Figure 9, the rising autocorrelation in consumption growth first drives the risk-free rate down as the autocorrelation rises toward zero, and thereafter drives it up as the autocorrelation becomes more positive. Meanwhile, the falling standard deviation should drive it up. Thus the effects work together except at the beginning of the sample period, when the consumption growth autocorrelation is negative and rising. But the reduction in mean consumption growth early in the sample clearly dominates, driving the risk-free rate down. From about the 1940-sample on, increasing mean and autocorrelation together with falling standard deviation all work together to push the model risk-free rate up (Figure 7b).

The dramatically counterfactual predictions of the habit model (Figure 7) in response to the changing consumption growth moments in the data naturally lead to the question of whether alternative parameterizations of the model could improve its performance. The answer is no, since given the consumption moments in the data, any *fixed* parameterization of the model will predict, counterfactually, a declining equity premium and rising risk-free rate (as in Figures 7a and 7b). For instance, $\beta = 0.96$, $\delta =$

-0.72 and σ = 0.53 match the equity premium and risk-free rate for the 1889-1978 sample used by Mehra-Prescott and Constantinides. Holding this parameterization fixed, the model makes counterfactual predictions for the evolution of the equity premium and risk-free rate similar to those in Figure 7. In fact, for two-thirds of the 40-year subsamples, it is possible to find apparently reasonable (δ, σ) combinations (holding β = 0.96) that will produce the observed equity premium and risk-free rate.[11] Choosing any one of these parameter combinations yields results that are qualitatively and quantitatively similar to Figure 7. Figure 10 provides some intuition about what is "apparently reasonable."

The remarkable feature of the habit model revealed by Figure 10 is that small values of σ, combined with moderate values of δ, can produce enormous values of the equity premium and small risk-free rates. In addition, for fixed δ around -0.5, very small increases in σ lead to huge increases in the equity premium and reductions in the risk-free rate. The figure thus reveals another dimension along which the habit model is extraordinarily sensitive to its parameterization. (Note that the time-separable model, i.e., δ = 0.0, does not share this sensitivity to changes in σ.)

Figure 10: Equity Premium and Risk-Free Rate for Different Preference Parameters

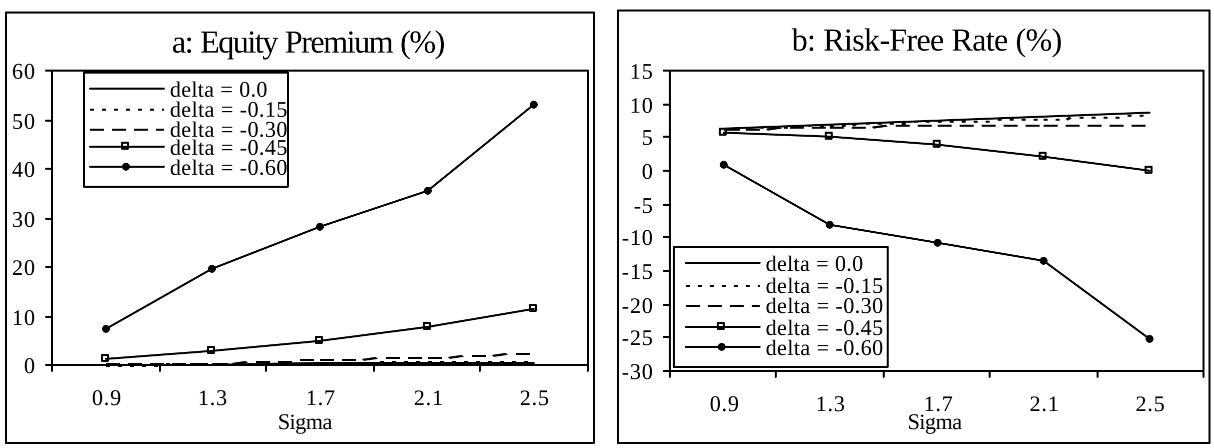

V. More General Forms of Habit Formation

In this section, we show that our results are robust to significant alterations in the asset-pricing model. In section V.1 below, we use the preferences described in (5), but allow for more than 1 lag in consumption to affect the utility in a given period. In V.2, we examine a model with *external* habit preferences.

[11] If one uses post-WWII *quarterly* U.S. data, the values of δ and σ that reproduce the equity premium and risk-free rate are β = 0.98, δ = -0.63 and σ = 6.3; this parameterization makes the model *extremely* sensitive to even the small

V.1 Multi-lag habit

While the 1-lag version of habit has enjoyed widespread use, a number of studies using habit formation preferences have included more than 1 lag of consumption in the habit stock. To capture consumption complementarities at longer horizons we use a geometrically decreasing lag structure for habit that was proposed and analyzed in Heaton (1995).[12] This version allows for a potentially infinite number of lags of past consumption to enter the current period utility function, and can be written as:

(8) $$E_o \sum_{t=0}^{\infty} \beta^t \frac{[(1+\delta(L))c_t]^{1-\sigma}}{1-\sigma} \text{ where } \delta(L) = -[|\delta|L^1 + |\delta|^2 L^2 + |\delta|^3 L^3 + \ldots].$$

Using the asset pricing procedure described in the previous section along with the estimated consumption growth process, we find that parameters $|\delta| = 0.41$, $\sigma = 0.7$ and $\beta = 0.975$ match the equity premium and risk free rate from 1889-1992.

Figure 11 displays the spectral utility function for this preference form. As with the one-lag habit model, the spectral utility function is downward sloping, indicating an aversion to high frequency volatility. The reason the spectral utility functions are similar is that with the geometrically decreasing polynomial the agent now has a greater dislike of lower frequency volatility, but the dislike of high frequency volatility remains.[13]

As suggested by the shape of the spectral utility function, the asset pricing implications we found in section III.2 hold as well. Holding overall volatility constant, as in Table 1, the consumption growth process with an autocorrelation of -0.3 yields an equity premium of 18.13% and the consumption growth process with autocorrelation 0.3 yields an equity premium of 2.59%—a 1550 basis point difference. Holding constant consumption volatility with cycles between 2 and 3 years results in a higher equity premium for the consumption growth process with serial correlation of 0.3. The equity premium is 2.66% with the negatively autocorrelated consumption growth process and 12.32% with the positively autocorrelated consumption growth process. This is due to two factors: first, overall consumption volatility has increased by a factor of 3 in making consumption more persistent; second, the geometric habit agent has much greater aversion for the now more pronounced 3-6 year fluctuations (compare the spectral utility in Figure 11 to the 1-lag habit spectral utility in Figure 1). If we hold constant the volatility between 2 and 6 years, the equity premium in the positively correlated consumption growth process case is only 130 basis points higher than that in the negatively autocorrelated consumption growth case. The result is very striking because the volatility between 2 and 6 years is only 2.4 percent of the overall

---

changes in the consumption process observed in the data for this period. As a result, the post-WWII version of Figure 7 is equally dramatic.

[12] Heaton also allows for durability.

[13] Otrok (2001) provides a more thorough explanation and illustration of the interaction between the lag polynomial in the utility function and the dislike of different types of consumption volatility.

volatility. Holding low frequency volatility constant, as in Table 3, the negatively autocorrelated consumption growth process yields an equity premium that is 1750 basis points greater. Clearly, the multi-lag habit specification is unlikely to alter the results on rolling subsamples in section IV.2 in a significant way.

Figure 11: Spectral Utility Function for Multi-lag Internal Habit

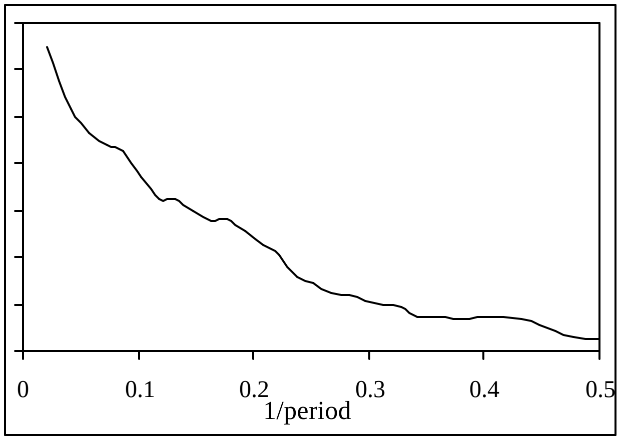

V.2 External habit

Campbell and Cochrane (1999, henceforth CC) develop an external version of habit formation commonly referred to as 'catching up with the Joneses'.[14] With this sort of habit formation the representative agent cares about his consumption relative to that of others. The preferences in Heaton (1995) and in CC both have the property that past consumption (potentially infinitely many periods in the past) affects the current habit stock. That is, habit evolves slowly. However, in CC (1999), since the habit stock is external, the representative agent does not take into account the effect of his consumption choice on the evolution of the habit stock.

The preference specification that CC employ is a non-linear function designed to match not only the mean equity premium, but also other moments of asset returns. We follow Ljungqvist and Uhlig (2000) and work with a linearized version of CC's preferences. Specifically, we use the linearized version presented in equation 2.3 of Campbell and Cochrane (1995). In what follows the linearization is not important.[15]

The period utility is given by:

[14] See also Abel (1990, 1999).

[15] The nonlinearity in the CC preferences is not needed to match the mean risk free rate and equity premium. Furthermore, the linearized version retains the important property that the habit stock adjusts slowly to changes in consumption.

$$(9) \qquad \frac{(c_t - X_t)^{1-\sigma}}{1-\sigma},$$

where X_t denotes the external habit stock. As in CC (1995) the habit stock evolves according to:

$$(10) \qquad \ln(X_t) = a + \phi \ln(X_{t-1}) + (1-\phi)\ln(C_{t-1}),$$

where C denotes aggregate consumption. In equilibrium, c = C. With ϕ = 0.83, a = -0.06, σ = 1.4 and β = 0.99, we are able to match the mean risk free rate and equity premium.

Figure 12 depicts the spectral utility function for the external habit formation model. This agent, like the agent with internal habit formation, has a relative preference for low frequency volatility over high frequency volatility. Also, Figure 13 shows that the external habit model retains the property that increasing the variance of consumption volatility equally at all frequencies has a disproportionate effect at higher frequencies. However, in contrast to Figure 2b, Figure 13 shows that the relative dislike of high frequency volatility is spread over a larger range of frequencies.

Figure 12: Spectral Utility Function for External Habit

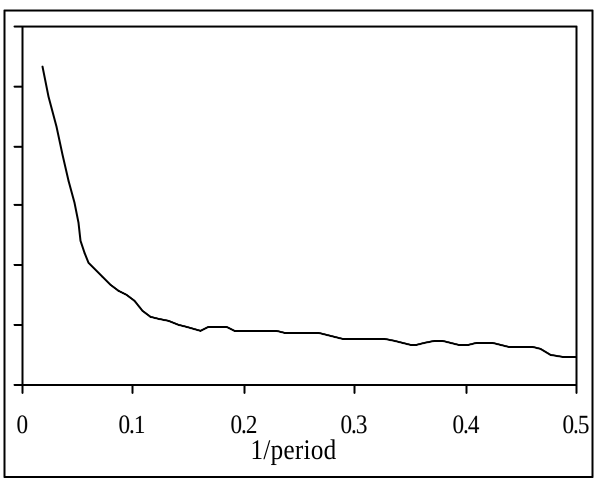

Figure 13: Spectral Utility Function for Different Consumption Volatilities

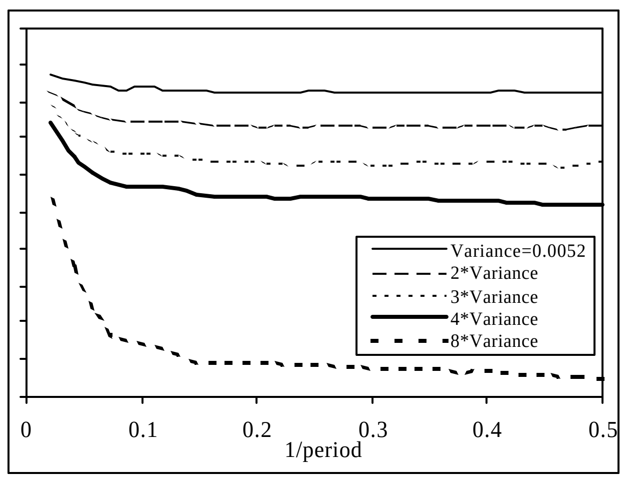

As with the one-lag habit model of Section III.2, the implications of changing the spectral characteristics of consumption for equity premium are dramatic. For instance, holding constant the overall volatility in consumption growth, in moving from growth autocorrelation 0.3 to –0.3, the equity premium rises by more than 750 basis points. Holding constant the volatility at higher frequencies (here defined as 2-7 period cycles), the equity premium rises, but only by 50 basis points. Note that as we allow the habit stock to include more of the history of consumption we include cycles of longer length (cycles of length 2-3 years for 1-lag internal habit formation and cycles of length 2-6 years for multi-lag internal habit formation). However it is important to note that the volatility in the range of 2-7 period cycles represents only 6 percent of the overall volatility.

Figure 14 below shows that the external habit model, like the internal habit model, is sensitive to the changes in consumption growth moments over time. The equity premium and risk free rate in the external habit model moves in a fashion qualitatively similar to the case of the 1-lag internal habit model (Figure 7). Quantitatively the model is not quite as sensitive, but the movements remain dramatic: comparing the early subsamples (ending in 1932 or 1933) with the latest ones, the equity premium exhibits a 1600 basis-point decline with the external habit model; the decline in the internal habit model was over 2500 basis points.

Figure 14: The Equity premium and Risk Free Rate Over Time

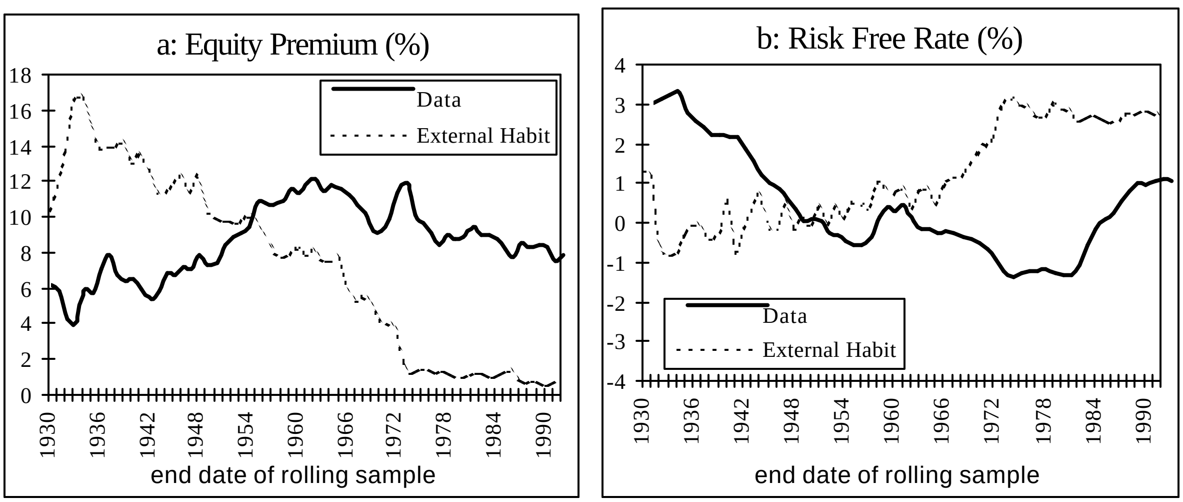

VI. Conclusion

Agents with habit preferences care not only about the volatility of the consumption process they face, but also about the temporal distribution of the variance. This phenomenon was illustrated using spectral utility functions, which decompose preferences for volatility by frequency (Figures 1, 2, 11, 12, and 13). Given overall volatility, the habit agents prefer more serial correlation to less. One can thus derive "habit indifference" curves depicting how these agents are indifferent between high overall volatility concentrated at low frequencies and low overall volatility concentrated at high frequencies. That is, the habit agents can maintain the same expected utility while taking on greater volatility, provided the volatility comes at lower frequencies. This illustrates how the conventional intuition based on mean-variance tradeoffs must be modified in a world with time nonseparabilities.

Our spectral utility functions are also useful for understanding the role of time-non-separable preferences in explaining other stylized facts. For example, in the business cycle literature Kydland and

Prescott (1982) argue that non-separabilities in preferences are necessary to explain the smoothness of wages relative to hours worked. Barro and King (1984) develop a model with time non-separable preferences that can rationalize the observed procyclicality of hours. Backus, Gregory and Telmer (1993) use habit preferences to try to explain the predictable returns from exchange rate speculation.

We have used spectral utility here to help isolate the route through which the habit model 'solves' the equity premium puzzle: it is founded on the heightened aversion habit agents have to high-frequency consumption fluctuations. Indeed, the habit model delivers the equity premium by exhibiting extraordinary sensitivity to high-frequency fluctuations. Because of this sensitivity, as the volatility and serial correlation properties of U.S. consumption have changed during the last 100 years, the model makes dramatically counterfactual predictions of the time path of the equity premium and the risk-free rate. Whether habit, as usually formulated, constitutes a resolution of the equity premium puzzle is therefore in question.

References

Abel, Andrew B., (1990), "Asset Pricing Under Habit Formation and Catching Up With The Joneses," *American Economic Review* 80:38-42.

Abel, Andrew B., (1999), "Risk Premia and Term Premia in General Equilibrium," *Journal of Monetary Economics* 43:3-34.

Backus, David K., Allan W. Gregory, and Chris I. Telmer, (1993), "Accounting for Forward Rates in Markets for Foreign Currency," *Journal of Finance* 48:1887-1908.

Bansal, Ravi, Wilbur John Coleman, (1996) "A Monetary Explanation of the Equity Premium, Term Premium, and Risk-Free Rate Puzzles," *Journal of Political Economy* 104:1135-71.

Barro, Robert, and Robert G. King, (1984), "Time-Separable Preferences and Intertemporal Substitution Models of Business Cycles," *Quarterly Journal of Economics* 99:817-839.

Baxter, Marianne, and Robert G. King, (1999), "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series," *Review of Economics and Statistics* 81: 575-593.

Boldrin, Michele, Lawrence J. Christiano, and Jonas D. M. Fisher, (1997), "Habit Persistence and Asset Returns in An Exchange Economy," *Macroeconomic Dynamics* 1:312-332.

Campbell, J. Y., and J. H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy* 107, 205-251.

Campbell, J. Y., and J. H. Cochrane, 1995, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," NBER working paper 4995.

Christiano, L.H., and T. J. Fitzgerald, 1999, "The Band Pass Filter," Working paper no. 7257, National Bureau of Economic Research, Cambridge.

Constantinides, George C., (1990), "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy* 98:519-43.

Golob, John, (1992), "A Regime Shift in Measured Per Capita Consumption, with Implications for Asset Prices and Returns," Federal Reserve Bank of Kansas City, research working paper RWP 92-06.

Heaton, John, (1995), "An Empirical Investigation of Asset Pricing with Temporally Dependent Preference Specifications," *Econometrica* 63:681-717.

Jermann, Urban J., (1998), "Asset Pricing in Production Economies," *Journal of Monetary Economics* 41:257-276.

Judd, K. (1998), *Numerical Methods in Economics*, MIT press, Cambridge.

Kocherlakota, Narayana, (1996), "The Equity Premium: It's Still a Puzzle," *Journal of Economic Literature* 34:42-71.

Kydland, Finn, and Edward C. Prescott, 1982, "Time To Build and Aggregate Fluctuations," *Econometrica* 50:1345-1370.

Ljungqvist Lars, and Harald Uhlig, (2000), "Tax Policy and Aggregate Demand Management Under Catching Up with the Joneses," *American Economic Review* 90:356-366.

Lucas, Robert E. Jr., (1978), "Asset Prices in an Exchange Economy," *Econometrica* 46:1429-1445.

Mehra, Rajnish and Edward C. Prescott, (1985), "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15:145-161.

Otrok, Christopher, (2001), "Spectral Welfare Cost Functions," *International Economic Review* 42:345-367.

Shiller, Robert J., (1989), *Market Volatility*, MIT Press, Cambridge MA.

Taub, Bart, (1989), "Dynamic Consistency of Insurance Contracts Under Enforcement by Exclusion," *Journal of Economic Dynamics and Control* 13:93-112.

Weil, Philippe, (1989), The Equity Premium Puzzle and The Risk-Free Rate Puzzle," *Journal of Monetary Economics* 24:401-421.

Whiteman, Charles H. (1985), "Spectral Utility, Wiener-Hopf Techniques, and Rational Expectations," *Journal of Economic Dynamics and Control* 9:225-240.

Whiteman, Charles H. (1986), "Analytical Policy Design Under Rational Expectations," *Econometrica* 54: 1387-1405.

Appendix: Calculating the Equity Premium and the Risk Free Rate

This Appendix provides details on calculating asset prices for the three forms of preferences we use in the paper. Section A.1 describes a general algorithm for calculating prices. Section A.2 provides details on the intertemporal marginal rates of substitution for the preference forms used in the paper.

A.1 General Procedure for Calculating Asset Prices

We follow the method described in Judd (1998, pp. 574-576) to calculate asset prices. A minor technical difference between our computations and Judd's is that Judd uses consumption levels, while we work with consumption growth rates. However, the algorithm and motivation are very similar. We start with the agent's first-order conditions for choosing the optimal consumption and shareholding sequences:

(A1) $$u'(c_t)p_t = \beta E_t u'(c_{t+1})(p_{t+1} + d_{t+1}), \qquad t \geq 0.$$

Divide through by $u'(c_t)$ to obtain:

(A2) $$p_t = \beta E_t \frac{u'(c_{t+1})}{u'(c_t)}(p_{t+1} + c_{t+1}).$$

Note that we have substituted in the equilibrium condition that consumption equals dividends. Since prices are homogenous of degree one in consumption we can write $p_t = w_t c_t$, where w_t is the 'stationary' price. Substituting in for p_t we get:

(A3) $$w_t c_t = \beta E_t \frac{u'(c_{t+1})}{u'(c_t)}(w_{t+1}c_{t+1} + c_{t+1}).$$

Next, divide through by c_t to obtain:

(A4) $$w_t = \beta E_t m(\partial c_{t+1})(w_{t+1} + 1)\partial c_{t+1},$$

where ∂c_{t+1} is the gross growth rate of consumption and $m(\partial c_{t+1})$ denotes that the agent's intertemporal marginal rate of substitution (IMRS) is written using consumption growth rates rather than consumption levels. Section A.2 of this appendix provides details on the IMRSs for the 3 types of preferences that we use. We also have similar expressions for w_{t+1}, w_{t+2}, etc. By recursively substituting in for w_{t+1}, w_{t+2}, …, in (A4) we obtain:

(A5) $$w_t = \left[E_t \sum_{j=1}^{\infty} \beta^j \left(\prod_{i=1}^{j} m(\partial c_{t+i})\partial c_{t+i} \right) \right] + E_t \beta^{\infty} \left(\prod_{i=1}^{\infty} m(\partial c_{t+i})\partial c_{t+i} \right) w_{\infty}$$

Since β is less than one the last term goes to 0. Returns are given by:

(A6) $$R_t = \frac{w_{t+1}c_{t+1} + c_{t+1}}{w_t c_t} = \frac{[w_{t+1} + 1]}{w_t}\frac{c_{t+1}}{c_t}$$

To calculate the expectation in (A5) we use a simulation procedure. First we assume that consumption growth follows an AR(1) process:

(A7) $\partial c_{t+1} = \alpha + \rho \Delta c_t + \varepsilon_t$, where ε_t is $N(0, \sigma_\varepsilon^2)$.

We estimate the parameters α, ρ and σ_ε^2 with OLS. The procedure for calculating w_t is as follows:

1) Create a discrete grid of values for ∂c_{t+1}. (In our simulations we use 30 grid points that cover a range from 3.5 standard deviations above and to 3.5 standard deviations below the mean of consumption growth.)
2) Simulate a long time series for ∂c_{t+j}, j = 2,3,… using the estimated parameters of the consumption growth process and using the first grid point for ∂c_{t+1} as an initial condition. (In our simulations we use a time series of length 200.)
3) With the consumption growth time series from 2) and a parametric form for utility evaluate equation (A5).
4) Repeat step 3) many times, and the average of the simulations is the price of the risky asset when consumption growth takes on the value of the 1st grid point. (In our simulations we draw 1500 time series.)
5) Repeat Steps 2)-4) for each of the grid points.
6) To get a pricing function project the vector of asset prices (w_t) from step 5 onto the consumption growth grid points and a constant.
7) To calculate a return time series, simulate a long time series for consumption growth. (Again, our simulations are of length 200.)
8) Using the regression coefficients from step 6) calculate a time series of prices. Next, using realized consumption growth and the price time series calculate the return time series using equation (A6).
9) Repeat steps 7) and 8) many times and average to calculate expected returns. (We use 1500 simulations.)

Judd (1998) provides details on the issues involved with this calculation. One important issue is the length of the time series used to approximate the infinite sum in (A5). Given our value of β, a time series of length 200 provides a good approximation since $0.955\wedge^{200}$ is close to zero. That is, additional terms have practically no impact on the sum in (A5). A second issue is the number of simulations used to calculate the expectation. We experimented with increasing the number of simulations to 3000, and the results were unchanged. A third issue is whether or not a linear function provides a good approximation to the true pricing function. We plotted the simulated prices for each grid point, along with the fitted values from the regression in step 6) and the pricing function in the time separable case is well approximated by a linear function. In the habit case the simulated values appear to have a small degree of non-linearity, though the fitted values appear close.

A.2 Intertemporal Marginal Rates of Substitution

This section provides details on calculating the IMRS for the models that we study.

Time-separable preferences

The first is CRRA time-separable preferences of the form:

(A8) $$E_o \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma}, \qquad \sigma > 0.$$

The corresponding IMRS, written in terms of consumption growth rates is:

(A9) $$m(\partial c_{t+1}) = \beta E_t \frac{c_{t+1}^{-\sigma}}{c_t^{-\sigma}} = \beta E_t (\partial c_{t+1})^{-\sigma}.$$

1-lag habit

Our first version of habit-formation preferences, which contains 1 lag of consumption, is of the form:

(A10) $$E_o \sum_{t=0}^{\infty} \beta^t \frac{[c_t + \delta c_{t-1}]^{1-\sigma}}{1-\sigma}, \delta < 0.$$

The corresponding IMRS is

(A11) $$\frac{u'(c_{t+1})}{u'(c_t)} = \beta E_t \frac{(c_{t+1} + \delta c_t)^{-\sigma} + \beta\delta(c_{t+2} + \delta c_{t+1})^{-\sigma}}{(c_t + \delta c_{t-1})^{-\sigma} + \beta\delta(c_{t+1} + \delta c_t)^{-\sigma}},$$

which can be written in terms of current and future consumption growth rates as follows

(A12) $$m(\partial c_{t+2}, \partial c_{t+1}, \partial c_t) = \beta E_t \frac{(\partial c_{t+1} + \delta)^{-\sigma} + \beta\delta[\partial c_{t+1}]^{-\sigma}(\partial c_{t+2} + \delta)^{-\sigma}}{(1 + \delta[\partial c_t]^{-1})^{-\sigma} + \beta\delta(\partial c_{t+1} + \delta)^{-\sigma}}.$$

Multi-lag habit

The second version of habit that we use allows for an infinite number of lags of past consumption to enter the period utility function:

(A13) $$E_o \sum_{t=0}^{\infty} \beta^t \frac{[(1+\delta(L))c_t]^{1-\sigma}}{1-\sigma}$$ where $\delta(L) = -[|\delta|L^1 + |\delta|^2 L^2 + |\delta|^3 L^3 + \ldots]$.

The geometrically decreasing lag structure for habit was proposed and analyzed in Heaton (1995). The IMRS for these preferences is:

$$\beta E_t \frac{\left(c_{t+1}-\sum_{i=1}^{\infty}\delta^i c_{t+1-i}\right)^{-\sigma}-\sum_{j=1}^{\infty}\beta^j\delta^j\left(c_{t+1+j}-\sum_{i=1}^{\infty}\delta^i c_{t+1+j-i}\right)^{-\sigma}}{\left(c_{t}-\sum_{i=1}^{\infty}\delta^i c_{t-i}\right)^{-\sigma}-\sum_{j=1}^{\infty}\beta^j\delta^j\left(c_{t+j}-\sum_{i=1}^{\infty}\delta^i c_{t+j-i}\right)^{-\sigma}},$$

or in terms of growth rates: $\beta E_t(\text{num}/\text{den})$ where:

$$\text{num} = \left[\begin{array}{l}\left(\partial c_{t+1}-\delta-\sum_{i=2}^{\infty}\delta^i\left(\prod_{j=1}^{i-1}\partial c_{t-j+1}\right)^{-1}\right)^{-\sigma} \\ \quad -\sum_{k=1}^{\infty}\beta^k\delta^k\left(\prod_{j=1}^{k+1}\partial c_{t+j}-\left(\sum_{i=1}^{k}\delta^i\left(\prod_{j=1}^{k+1-i}\partial c_{t+j}\right)\right)-\delta^{k+1}-\left(\sum_{i=k+2}^{\infty}\delta^i\left(\prod_{j=0}^{i-k-2}\partial c_{t-j}\right)^{-1}\right)\right)^{-\sigma}\end{array}\right]$$

$$\text{den} = \left[\begin{array}{l}\left(1-\sum_{i=1}^{\infty}\delta^i\left(\prod_{j=0}^{i-1}\partial c_{t-j}\right)^{-1}\right)^{-\sigma}-\beta\delta\left(\partial c_{t+1}-\delta-\sum_{i=2}^{\infty}\delta^i\left(\prod_{j=0}^{i-2}\partial c_{t-j}\right)^{-1}\right)^{-\sigma} \\ \quad -\sum_{k=2}^{\infty}\beta^k\delta^k\left(\prod_{j=1}^{k}\partial c_{t+j}-\left(\sum_{i=1}^{k-1}\delta^i\left(\prod_{j=1}^{k-i}\partial c_{t+j}\right)\right)-\delta^{k}-\left(\sum_{i=k+1}^{\infty}\delta^i\left(\prod_{j=0}^{i-k-1}\partial c_{t-j}\right)^{-1}\right)\right)^{-\sigma}\end{array}\right]$$

In the computations we approximate the infinite sums by truncating the summation at 10 lags. Given our estimates of β and δ, after about 10 lags the additional items in the sequence are very close to 0. Additionally, while we use the same time series of length 200 for asset pricing, we simulate a longer time series so that we have sufficient leads and lags to have an asset price for each of the 200 periods.

The simulation procedure decribed in A.1 will not work with the Heaton type preferences since lagged consumption growth enters the IMRS along with current and future consumption growth. For these preferences a slower procedure was used. First, a consumption time series is drawn. Second, at each date consumption is simulated forward many times to evaluate the expectation and to get a price for that date. With the price time series and consumption growth (dividend growth), one realization of the return time series can be calculated. Third, to calculate expected returns we have to repeat steps 1 and 2 many times.

External habit

The agents IMRS is given by:

(A14) $$m_{t+1}=\frac{\beta(c_{t+1}-X_{t+1})^{-\sigma}}{(c_t-X_t)^{-\sigma}}.$$

In order to obtain an expression for the IMRS in terms past and future growth rates of consumption (which is useful for computing asset prices), first re-write (A14) as:

(A15) $$m_{t+1} = \left(\frac{\frac{c_{t+1}}{c_t} - \frac{X_{t+1}}{c_t}}{1 - \frac{X_t}{c_t}} \right)^{-\sigma} = \left(\frac{\frac{c_{t+1}}{c_t}\left(1 - \frac{X_{t+1}}{c_{t+1}}\right)}{1 - \frac{X_t}{c_t}} \right)^{-\sigma} .$$

We now have an IMRS in terms of the consumption growth rate and the ratio of the habit stock to consumption. To get an expression for the habit/consumption ratio as a function of consumption growth, subtract $\ln(c_{t+1})$ from both sides of (10) and collect terms to get:

(A16) $$\ln(X_t) - \ln(c_t) = a + \phi(\ln(X_{t-1}) - \ln(c_{t-1})) + \ln(c_{t-1}) - \ln(c_t) \text{ , or}$$

(A17) $$\ln(\frac{X_t}{c_t}) = a + \phi \ln(\frac{X_{t-1}}{c_{t-1}}) - \ln(\frac{c_t}{c_{t-1}}) .$$

Repeatedly substituting for the lagged habit/consumption ratio yields:

(A18) $$\ln(\frac{X_t}{c_t}) = \sum_{j=0}^{\infty} \phi^j a - \sum_{j=0}^{\infty} \phi^j \ln(\frac{c_{t-j}}{c_{t-1-j}})$$

With this expression we can evaluate (A15).

To calculate the price of the two assets we follow the procedure described above for multi-lag internal habit preferences. There is one additional issue: as the level of consumption falls to the level of the habit stock there is potential for habit to exceed consumption, in which case the utility function is no longer well defined. We follow CC (1999) and impose an upper bound (0.95) on the ratio of habit to consumption (X_t/C_t). As in the simulations reported by CC, this limit rarely binds.