Consumption adjustment costs and the equity premium

Narayan Bulusu[a] and Javier Gómez Biscarri[b]

[a] *Bank of Canada*

[b] *Universitat Pompeu Fabra and Barcelona Graduate School of Economics*

January 15, 2012

Abstract

We show that two types of consumption risks are priced in the equity premium: the risk of aggregate consumption growth and that of changing the composition of the consumption basket, when goods have heterogenous costs of adjustment. Using the property that frequency of consumption adjustment is inversely related to adjustment costs, we identify low adjustment cost goods by the high-frequency component of nondurables and services consumption obtained using band pass filters. These data help resolve the equity premium puzzle, while simultaneously generating sufficient volatility of marginal utility to cross the lower volatility bound at all frequencies.

1 Introduction

The difficulties of the classical CCAPM in explaining the observed equity premium are driven by the low volatility of consumption growth, which requires infeasibly large values of the risk aversion parameter to satisfy the Hansen and Jagannathan (1991) (HJ) bound. Recent work in the frequency domain has suggested that a closer look at the cyclical components of the data and of the model-generated variables (i.e., the relative importance of the frequencies of fluctuation of different components in these variables) may increase our understanding of the empirical issues involved: Berkowitz (2001) and Otrok, Ravikumar, and Whiteman (2007) show that the CCAPM performs well at low frequencies - components of the equity premium that are slow-moving - whereas its failures are concentrated at the higher frequencies. In this paper, we propose that one way to improve the high-frequency performance of the CCAPM - and consequently, resolve the equity premium puzzle - is by considering the heterogeneity of adjustment costs.

The various components of nondurable goods and services[1] are treated as perfect substitutes in the traditional implementations of the CCAPM. There is at least one critical dimension along which they are not: the cost borne by the consumer to adjust consumption. It is substantially costlier to move house (rental services consumption) than to change brands of cereal.[2] An agent who suffers a wealth shock would attempt to adjust consumption of goods with lower adjustment costs before changing consumption of the high adjustment cost goods. In other words, goods with lower adjustment costs exhibit high-frequency changes, whereas those with higher adjustment costs are not adjusted at high frequencies. This has two results: first, the volatility of consumption growth is higher for components with lower adjustment costs; second, the high-frequency component of growth of the wealth portfolio is more correlated with the growth of consumption of the low adjustment cost goods than with the growth of the high adjustment cost goods.

Our agent is not just concerned about total consumption, but critically, also on the evolution of the ratio of the two composition of the consumption basket. Since the agent does innately prefer one of the two types of goods, a change in the ratio of consumption of the high to low adjustment cost goods implies that the adjustment costs 'bite'. When the two goods are not perfectly substitutable, this has an impact on her marginal utility.

The agent's adjustment of the two components of consumption depends on the nature and magnitude of the wealth shock. For shocks that the agent perceives as temporary, and especially for shocks of small magnitude, the agent finds it optimal to maintain the consumption level of the goods with high adjustment costs, but compensate by changing consumption of goods with low adjustment costs. Thus, for short horizons, the high-frequency component of consumption drives the marginal utility. On the other hand, for shocks that are more long-run in nature, especially for those with large magnitudes, the agent is forced to adjust consumption of goods with high adjustment costs. Thus, for long horizons, the low-frequency component of consumption is the key driver of the marginal utility. The CCAPM aggregates consumption at all frequencies, and since the low-frequency dominates the total consumption basket,[3] the agent's marginal utility is approximated by that of the consumption of the low-frequency component. This explains why the CCAPM does well at low frequencies but fails to explain the high-frequency movement of equity returns: a very large proportion of the consumption basket is composed of consumption that does not change at high frequencies; aggregation of all non-

[1]Since durable goods are not used in this paper, we hereafter refer to nondurable goods and services simply as 'goods'.

[2]The costs of adjustment are not just fixed costs: moving house involves not just paying the cost of transportation of goods and the costs involved in searching for children's schools, but also changes in cost of transportation to the workplace and the like.

[3]In our estimates, we find that the high-frequency component of consumption forms just over 5% of total nondurable and services consumption.

durable and services consumption therefore obscures the high-frequency consumption changes which are key to explaining high-frequency movements of equity returns.

Monthly or quarterly data at sufficient granularity are unavailable to allow us to accurately classify consumption based on the costs consumers face in adjusting them. Marshall and Parekh (1999) show that in the presence of adjustment costs, the consumer changes consumption less frequently; in other words, adjustment costs lower the frequency of consumption adjustment. Further, they show that the 'range of inaction' increases with the magnitude of the adjustment cost: the larger the cost, the lesser the frequency of adjustment. We use this relationship between costs of adjustment and frequency of adjustment in obtaining the time series of the two types of consumption. We use band pass filters on the series of monthly nondurable and services consumption to isolate the component of consumption with cycle length less than one year, and use this as a proxy for the basket containing goods with low adjustment costs. Goods with high adjustment costs are then simply obtained by subtracting the filtered series from the original consumption series.[4]

Our results show that using this split of consumption helps resolve the equity premium puzzle: the estimate of risk aversion is 'more reasonable', and the expected risk-free rate is lower than that obtained using the CCAPM. The model estimation errors are also lower, primarily because the SDF is more volatile. Analysis in the frequency domain shows that the SDF satisfies the HJ bounds at all frequencies.

We check robustness by using alternate band pass filters and alternate definitions of what constitutes high-frequency consumption. We also perform a comparative analysis of the CCAPM and our modified version on other test portfolios. The results show that splitting consumption by adjustment costs improves model fit and also helps satisfy the HJ bounds at all frequencies. We also check that our proxy for goods with low adjustment costs is not identifying other consumption series - such the luxury goods factor in Aït-Sahalia, Parker, and Yogo (2004) nor the consumption of durable goods, shown to be important in Yogo (2006) - that have been used to explain the cross-section of equity returns.

Our paper is related to two strands of literature: studies that identify issues with the measurement or usage of consumption data in the implementation of the CCAPM;[5] and those that study the fit of the CCAPM in the frequency domain. Beginning with the observation that the observed consumption expenditure data are too smooth to explain the equity premium, Marshall and Parekh (1999) show that the presence of even marginal costs of adjusting consumption reduces volatility of aggregate consumption.[6] In contrast to their set-up, heterogeneity in adjustment costs is the key to our results. Yogo (2006) shows convincing evidence that the inclusion of durable goods consumption helps explain both the time series and cross section of stock returns through the amplifying effect on the marginal utility of the simultaneous fall in consumption of durable goods along with nondurable and services consumption in bad times. Aït-Sahalia, Parker, and Yogo (2004) propose, alternately, the inclusion of utility for luxury goods. Their argument relies on funding purchases of luxury goods from the market portfolio, but not depending on the market return to fund the basket of subsistence goods. We propose that neither the inclusion of durable goods, nor the assumptions regarding luxury goods and market participation, is necessary for increasing the volatility of the SDF; the same result can be achieved in a setting with a representative agent with CRRA utility by recognising the heterogeneity in adjustment

[4]We make no claim of this being an optimal identification strategy of the two types of consumption, only that our method is both reasonable and easy to implement.

[5]The literature on the resolution of the equity premium puzzle is too large to cite. Some notable contributions involve alternative agent preferences (e.g. Epstein and Zin (1991), Campbell and Cochrane (1999)), the postulation of long-run risks (e.g., Bansal and Yaron (2004)) or rare disasters (e.g., Rietz (1988)).

[6]Other responses to the 'problem' of low volatility of aggregate consumption growth have been to suggest that consumption growth changes over several periods in response to changes in wealth (Parker and Julliard (2005)), or arguing that the data is poorly measured, and that other proxies are better (e.g., Savov (2011)).

costs of nondurable and services consumption.

A less extensive literature analyses the performance of the CCAPM in the frequency domain, where there is consensus that the performance of not just the CCAPM, but also all the proposed alternatives involving alternate utility specifications, is poor at the higher but good at the lower frequencies. For example, Berkowitz (2001) transforms observed model error correlations at various lags into the empirically observed spectrum and compares with the flat spectrum that should be observed for the 'right' model (errors of a 'correct' pricing model are uncorrelated at all lags), and finds that the deviation is highest for high frequencies. Cogley (2001) finds the minimum error that needs to be added to the model-generated discount factor to explain the data, with the additional restriction that the total satisfy the HJ bound, and then analyses the spectrum of the error. Otrok, Ravikumar, and Whiteman (2007) improves on this by estimating an equivalent of the HJ bound at all frequencies, compares it to the spectrum of the model-generated SDF, and agrees with Cogley (2001) that the CCAPM is unable to account for high-frequency movements of the equity premium. We use the analysis proposed by Otrok, Ravikumar, and Whiteman (2007) to complement the analysis of the properties of the model we propose in the time domain with that in the time domain, and show that our model does well in both.

In section 2, we describe a simple model where consumption is split into two components and use it to derive the SDF. In section 3, we discuss why adjustment costs justify a possible split of the consumption in terms of its frequency components and describe our empirical strategy and estimation methods. In section 4 we present the estimation of our suggested two components of consumption and compare them to the traditional CCAPM. We end in section 5 with some concluding comments.

2 The setup

We choose a setting that is close to the traditional implementation of the CCAPM, but which allows us to analyse the impact of heterogeneity in costs of adjustment between goods. We write a simple two-good model, in which the total consumption is divided into two components. Accordingly, our representative investor has utility U for the consumption bundle C_t that is given by

$$U(C_t) = \frac{(C_t)^{1-\gamma}}{1-\gamma}$$

where the consumption C_t is the following function of the consumption with low and high adjustment costs, c_t^l and c_t^h, respectively

$$C_t = \left[\left(c_t^l\right)^{\frac{\epsilon-1}{\epsilon}} + \left(c_t^h\right)^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon}{\epsilon-1}}$$

This is a simplified version of the specification used in Piazzesi, Schneider, and Tuzel (2007), where we set $\omega = 1$.

ϵ is the intratemporal elasticity of substitution: a measure of how willing the agent is to substitute consumption of one of the two goods by the other at a given time t. The larger the value of ϵ, the more substitutable the goods are. Below the threshold value of unity, the two goods are complements, and as $\epsilon \to \infty$, the utility function reduces to the standard CRRA. γ is the intertemporal elasticity of substitution: a measure or how willing the agent is to substitute consumption at t for consumption at $t + 1$.

The agent chooses consumption of the two types of goods c_t^l and c_t^h and the number of 'shares' ξ_t invested in the global wealth portfolio at each date t, with the objective of maximising her expected two-period utility

$$E_t \left[U(c_t^l, c_t^h) + \beta U(c_{t+1}^l, c_{t+1}^h)\right]$$

where β is the subjective discount factor. The agent receives an exogenous endowment e_t at each t, which means that her choice of consumption and investment is limited by

$$\begin{aligned} c_t^l + p_t^h c_t^h &= e_t - p_t \xi_t \\ c_{t+1}^l + p_{t+1}^h c_{t+1}^h &= e_{t+1} + p_{t+1}\xi_{t+1} \end{aligned}$$

where p_t^h is the price of the second consumption good, p_t is the price of equity, both in terms of the consumption good with low costs of adjustment; and ξ_{t+1} is the number of shares of the wealth portfolio held at $t+1$.

The first order condition derived with respect to ξ_t gives the well-known asset pricing equation $p_t = E_t(M_{t+1}p_{t+1})$ where the stochastic discount factor is given by

$$M_{t+1} = \beta \left(\frac{c_{t+1}^l}{c_t^l}\right)^{-\gamma} \left[\frac{1 + \left(\frac{c_{t+1}^h}{c_{t+1}^l}\right)^{\frac{\epsilon-1}{\epsilon}}}{1 + \left(\frac{c_t^h}{c_t^l}\right)^{\frac{\epsilon-1}{\epsilon}}}\right]^{\frac{1-\epsilon\gamma}{\epsilon-1}} \tag{1}$$

2.1 Comparison with the CCAPM

As mentioned earlier, the agent's utility reduces to the standard CCAPM utility specification when the elasticity of intertemporal substitution $\epsilon \to \infty$.[7] Heterogeneity in adjustment costs 'forces' the agent to differently adjust the two types of consumption; this matters when the two goods are not perfectly substitutable.

In the presence of (some degree of) insubstitutability, the evolution of the ratio of consumption of the two goods is an additional determinant of marginal utility. We show below that when the growth rate of the high adjustment cost goods is greater than that of the low adjustment cost goods, the agent has a higher marginal utility than is recognised under the traditional CCAPM specification. Conversely, the agent's marginal utility is lower than that specified under CCAPM when the growth rate of the low adjustment cost goods is grater than that of the high adjustment cost goods. Further, the ratio M_{t+1}/M_{t+1}^{CCAPM} is larger in magnitude the larger the change in consumption ratio c^h/c^l between the two time periods. Aggregate consumption growth is not the only risk factor in our model; the risk of a change in the composition of consumption is also key in determining the agent's marginal utility.

However, even if the goods are not perfectly substitutable, if the growth rate of the two kinds of consumption is identical, the agent behaves 'as if' the heterogeneity of adjustment costs does not matter, and the marginal utility is equivalent to that under the CCAPM.

We provide some more details and an intuition behind why the ratio of the consumption of the two kinds of goods matters to the consumer. To perform this comparison, we first rewrite the SDF obtained from the CCAPM as

$$\begin{aligned} M_{t+1}^{CCAPM} &= \beta \left(\frac{c_{t+1}^l + c_{t+1}^h}{c_t^l + c_t^h}\right)^{-\gamma} \\ &= \beta \left(\frac{c_{t+1}^l}{c_t^l}\right)^{-\gamma} \left[\frac{1 + \frac{c_{t+1}^h}{c_{t+1}^l}}{1 + \frac{c_t^h}{c_t^l}}\right]^{-\gamma} \end{aligned} \tag{2}$$

[7]Consistent with this, Figure 1 shows that the ratio of the SDF in the two cases gets closer to one the greater the value of ϵ. See the discussion in this section for details of its construction.

The reason we write it this way is now clear: a comparison of the SDFs in the two cases is simply a comparison of the third terms of equations (1) and (2).

In the appendix, we show the mechanism by which the SDF of our model has higher volatility: for reasonable values of the parameters ϵ, γ, and the ratio of consumption of goods with high to those of low adjustment cost c_t^h/c_t^l,

$$M_{t+1} > M_{t+1}^{CCAPM} \quad \text{iff} \quad \left(\frac{c_{t+1}^h}{c_{t+1}^l}\right) > \left(\frac{c_t^h}{c_t^l}\right) \tag{3}$$

$$M_{t+1} < M_{t+1}^{CCAPM} \quad \text{iff} \quad \left(\frac{c_{t+1}^h}{c_{t+1}^l}\right) < \left(\frac{c_t^h}{c_t^l}\right) \tag{4}$$

$$M_{t+1} = M_{t+1}^{CCAPM} \quad \text{iff} \quad \left(\frac{c_{t+1}^h}{c_{t+1}^l}\right) = \left(\frac{c_t^h}{c_t^l}\right) \tag{5}$$

For some intuition behind these conditions, consider first the reaction of the agent to a negative shock that is not very large (especially when the shock is judged to be transient): she attempts to lower her consumption of the low adjustment cost good while maintaining the level of consumption of the high adjustment cost good. In other words, she chooses $c_{t+1}^h = c_t^h$ and $c_{t+1}^l < c_t^l$; this raises her marginal utility because adjustment costs 'force' her to change the consumption of the two goods differently from what she would have ideally liked to do. A small positive shock, especially when it is transient, similarly induces the agent to choose $c_{t+1}^h = c_t^h$ and $c_{t+1}^l > c_t^l$; this lowers her marginal utility when compared to the CCAPM because the pleasure of increasing consumption comes without the risks of increasing the level of the high adjustment cost goods.

Negative shocks too large to be completely absorbed by adjusting consumption of the low adjustment cost goods lead to a necessary adjustment of the high adjustment cost goods. Given that high adjustment cost goods on the average comprise about 95% of total consumption basket in our sample, a large shock will lead to the situation where $(c_{t+1}^h/c_t^h) < (c_{t+1}^l/c_t^l)$, thus reducing the marginal utility. This might seem counter-intuitive at first glance: a large negative wealth shock results in a less unsatisfied consumer when compared to the CCAPM. However, if the shock is sufficiently large for the agent for it to be optimal to pay the adjustment costs, the consumption of goods with low adjustment costs experiences a lower decrease, which has positive marginal utility. The agent's behaviour when faced with a large positive shock is not clear-cut, and depends on which of the two types of goods she chooses to increase relatively more of.

Given that the poor empirical fit of the CCAPM stems from the low volatility of its marginal utility, the higher volatility of marginal utility generated by our modification is a step in the right direction. While marginal utility under some conditions in our proposed model is lower than under the CCAPM, Figure 1 shows that the increase of the SDF due to the split in consumption we propose is far larger than the (magnitude of the) decrease; thus the net impact of our proposed model is to increase the agent's marginal utility, or in other words, to make stocks riskier. This naturally leads to a larger equity premium.

In the next section, we describe the empirical procedure we use to confirm the results in this section, that suggest that separating consumption by its cost of adjustment would help better explain the data. We begin by describing the data and then discuss the method we employ to obtain proxies of the series of high and low adjustment components of consumption.

3 Data description

The data we use are standard and their are properties well-known, so we simply point to their sources and do not dwell on the details.

Data on total US monthly nominal nondurable and services consumption are available on the website of the US Bureau of Economic Analysis beginning January 1959. Data on consumer price index of nondurable goods and services is available from the website of the US Bureau of Labor Statistics; these data are used to calculate the monthly series of real nondurable and services consumption. Annual data of the population of the US are available from the website of the US Census Bureau, and a linear interpolation is applied to estimate the population in each month; these data are then used to estimate the real per-capita consumption of nondurable goods and services.

Data on all the test assets - the equity premium, the risk-free rate and the 25 portfolios sorted by size and book-to-market ratio - are obtained from Professor Ken French's website. All the data are available monthly beginning January 1959 up to December 2010. We always use these data to construct data at other time horizons, when necessary.

3.1 High and low frequency components of consumption

The method we use to identify components of consumption with high and low adjustment costs relies on the results of Marshall and Parekh (1999). They solve for the optimal consumption of heterogenous agents in the presence of fixed costs of adjustment. While their primary object is to show that even very small adjustment costs are sufficient to reproduce the observed volatility of aggregate consumption, we use another result they generate: consumers choose not to change their consumption if the difference between the optimal consumption in the absence of adjustment costs and consumption in the previous period is less than a quantity that is increasing in the cost of adjustment. The range of inaction in which the agent chooses not to change consumption increases with adjustment costs; in other words, consumption with high costs is infrequently changed - or, has low-frequency. Goods with low consumption costs, on the other hand, are adjusted frequently, or, have high-frequency. We therefore proxy consumption with low adjustment cost by the high-frequency component of nondurable and services consumption, and label the remaining consumption with high adjustment costs. In what follows, for brevity, we use the terms high-frequency and low adjustment costs interchangeably.

Given the importance of the method of extracting the low and high frequency components of consumption to our empirical analysis, we first discuss the band pass filters we use for this purpose and then present some characteristics of the resulting data.

Filters that isolate components of time series at desired frequencies have been well-known since Hodrick and Prescott (1997) propose a method to isolate what they call the 'growth component' of a time series by minimising a function of the difference between the original series and the desired growth component series ('HP filter'). Baxter and King (1999) point out that the specification of the λ parameter required for using this filter is an undesirable property, and instead propose a simple and popular alternative. The Cramer representation theorem (see, for example, Brockwell and Davis (1987) for details) shows that any stationary time series can be expressed as the integral of orthogonal random periodic components. Inverting this, the random component at each period can be written as an infinite weighted sum of the original time series. Baxter and King (1999) propose an approximate filter by truncating the ideal filter's weights at the desired lag ('BK filter'). Christiano and Fitzgerald (2003) propose an improvement to the approximation by projecting the ideal filter onto the space of the original time series ('CF filter'). While unlike the Baxter and King (1999) filter, their optimal filter is asymmetric, implying that the number of leads and lags used to estimate the frequency component is not equal, they also propose a symmetric version of their filter. Their filter has the additional feature that it can also be used with nonstationary series.

We use the four alternate band pass filters on the series of monthly real nondurable and services consumption - defining high-frequency as consumption that changes with cycle length less than six, 12, 18 and 24 months, separately - to estimate consumption with low adjustment costs.[8] Consumption with high adjustment costs is then simply the difference between the original consumption series and the filtered high-frequency component. Why draw the limit at 24 months for low-frequency consumption? We want to capture consumption with low adjustment costs, which are prone to frequent adjustment. Consumption adjusted as infrequently as once every two years already must have sufficiently high adjustment costs to disqualify it being so classified, especially given that we use the time series of nondurable goods and services. Why choose 6 months as the minimum, and not go lower? The band pass filters we use cannot identify components with cycle lenghts of one month, given that we have monthly data, so that when the maximum cycle length is six months, the filter actually passes through components with cycle lengths 2-6 months. The point is to set the definition in such a way as to make the maximum distinction between the adjustment costs of the two resulting series. This as ultimately an empirical question, which is why we use alternate definitions.

Passing a time series through a band pass filter alters its properties; we therefore present summary statistics of the filtered consumption series in Table 1, which shows that the key properties (in the time domain) do not depend on the type of filter used. The low frequency component forms the overwhelming part of the basket of consumption, about 95% on the average, ranging from 90.7% to 96.5% in the sample. The low frequency component has a time trend; the fact that it is nonstationary can be observed from the autocorrelation at various lags: the correlation of consumption with a lag of one year is as high as 94%. On the other hand, the high-frequency component is stationary. However, the other properties of the high-frequency component depend on the definition of 'high frequency' used in the filter. The lower part of each panel, which has similar statistics for the growth rates of the two components in the sample, shows that the high-frequency component has a higher mean and higher variance than the low-frequency component. This is as expected, since as argued earlier, consumption with high adjustment costs has lower variance; further, given agents are wary of increasing the consumption of this component for fear of 'being stuck' with the kind of consumption that cannot be easily changed. The growth of the low-frequency component is an AR(1) process with autocorrelation close to 1, but is stationary, as can be seen from the autocorrelation at higher lags. The growth of the high-frequency component has a significant negative autocorrelation at the first and fourth (not shown) lags, and significantly positive at the third.

Comparing the panels in Table 1, we see that the time series properties of the components of consumption obtained using the various filters are not very dissimilar. To add to this impression, see Table 2 which shows the correlation of the time series of high-frequency consumption obtained using the four filters discussed earlier. The two definitions of high-frequency used here are less than six months, and less than one year. The correlation between the consumption series using the same definition of high-frequency is close to 1. Correlation between series using different definitions of high-frequency is also high, between 75% and 80%. Despite this evidence that the series obtained using alternate filters are not materially very different, we conduct empirical analyses in the next section with all the alternatives.

Figure 2 contains preliminary evidence that we would expect to observe: high-frequency consumption drives the marginal utility - and consequently, the equity premium - at short horizons, while low-frequency consumption is the primary driver of the equity premium at longer horizons. We plot

[8]The code we use has been made available at the website of the Federal Reserve Bank of Atlanta.

the time series of equity premium and the growth of the two components of consumption, for various horizons. In the first row, for example, growth of consumption is calculated month-on-month, and the equity premium is the return on the equity over one month less the risk-free rate over one month. The high frequency of consumption is obtained by the asymmetric CF filter for consumption with cycle length less than one year. In the last row, the monthly consumption series is aggregated to calculate consumption for each year of the sample, and yearly growth rates are calculated. The equity premium is calculated as the product of the monthly gross returns on equity less the risk-free rate over a year. The graphs on the left side plot the equity premium and the growth rate of the high-frequency component of consumption and while those on the right side plot the equity premium and the growth rate of the low-frequency component. The correlation between the series in each graph is shown in the top-left corner of each plot. It can be seen that the correlation between the equity premium and the high-frequency component decreases with horizon while the correlation with the low-frequency component increases with the horizon.

In the next section, we show that considering splitting consumption into its high- and low-frequency consumption helps resolve the equity premium puzzle. We also test the performance of our model on the cross-section of stocks, represented by the 25 Fama-French portfolios sorted by size and book-to-market. Further, our low-frequency component is not highly correlated with other modifications to consumption proposed in the literature.

4 Empirical analysis

Our primary objective in proposing that we recognise heterogeneity in consumption costs is in resolving the equity premium puzzle, namely, to generate the empirically observed equity premium and risk-free rate, without recourse to a high coefficient of risk aversion. For this purpose, we use the monthly CRSP value-weighted portfolio and the return on the one-month T-Bill as the test assets. Since we test the conditional version of the CCAPM, we choose the lagged value of the CRSP market return and unity as the two instruments. Our first step is to estimate the parameters of the SDF in (1) and (2) in this set-up. The former involves estimation of three, and the latter, two, parameters. With two test assets and two instruments, this leaves one and two degrees of freedom for the parameter estimates, in the former and latter cases, respectively. Estimation is by a standard two-step GMM using the identity matrix in the first step; the optimal weighting matrix is constructed using 12 lags and Bartlett weights.

We complement this analysis by recent work that allows model diagnosis in the frequency domain. Otrok, Ravikumar, and Whiteman (2007) propose a method to check whether the SDF generated by the model has a volatility that is greater than the lower bound at every frequency. Hansen and Jagannathan (1991) estimate the lower bound by projecting the SDF onto the space of returns; they extend this by projecting the SDF onto the space of all past, current, and future returns. This allows the calculation correlations of the SDF at various leads and lags, whose Fourier transform gives the lower bound in the frequency domain. Comparing the spectrum of the model SDF and the lower bound helps identify the set of frequencies ω where the volatility of the model SDF for a frequency ω falls below its lower bound at the same frequency ω; in other words, the frequencies where the model performs poorly.

The results of the parameter estimation using the equity premium and the return on the one-month T-Bill are presented in Table 3. The top panel contains results of the CCAPM, which is in line with what we would expect: the risk aversion parameter is a rather high 16.7, and the stochastic discount factor has a very low volatility. The model is rejected, as can be seen from the p value of the null hypothesis that the error $\delta = 0$. Each of the subsequent four panels shows the estimation of the parameters of M_{t+1}, where the high-frequency component is estimated using the alternate filters discussed in the previous section. Within each panel, results are also presented for different definitions of the high-frequency component. Our model has errors that are lower than that of the CCAPM, irrespective of the filtering method or the split of consumption (errors are between 9% and 15% lower for the various alternatives). The risk-free rates produced by our model are also lower in all cases; in the majority of cases, the expected one period risk-free rate is less than two-thirds of what one would expect under the CCAPM. Next, the model SDF more than 30 times higher in all cases, and upto 60 times higher in some cases. Finally, the estimate of γ is lower in for all filtering methods, except for when the high-frequency component of consumption has frequency less than six months. We believe this could be due to the fact, pointed out earlier, that the band pass filters are unable to identify components that move at the very highest frequency, with cycle length one month. In all, the results show the ability of our modification to the standard CCAPM in resolving the equity premium puzzle. Despite this superior performance, however, the model is also resoundingly rejected by data: the p-values that the model errors are zero are of the order of 10^{-9}. This hurdle, however, has proved too high for nearly all models involving consumption data, and ours merely joins the back of the queue of an illustrious list of studies.

Figure 3, which plots the realisations of the SDF of the two models for case in which the Asymmetric CF filter is used for cycle length below one year, provides extra perspective on the greater volatility of the M_{t+1}. The deviations of M_{t+1} over 1 are noticeably greater in magnitude than its deviations below zero, when compared to M_{t+1}^{CCAPM}. The agent's degree of insatiation is larger, and she is more often insatiated, when compared to the agent with the CCAPM preferences. This is precisely what makes stocks command a greater premium, and helps resolve the equity premium. The question whether this increased volatility translates into better frequency-domain performance is answered by comparing Figures 4 and 5. The former figure plots the spectrum of M_{t+1}^{CCAPM} and the lower bound estimated using the procedure described in Otrok, Ravikumar, and Whiteman (2007), and the latter plots the spectra of M_{t+1} and their lower bound using the asymmetric CF filter for various definitions of high frequency. The former plot confirms results from other studies of the CCAPM in the frequency domain: it is unable to explain the high-frequency movements in the equity premium because the volatility of the high-frequency component of M_{t+1}^{CCAPM} is lower than the lower bound mandated by the data. Thus, CCAPM fails at high frequencies. While we see that the volatility of its low-frequency component is larger than the bound, its interpretation is merely that it satisifies the bound, and other model diagnostics must be used to check goodness of fit. This analysis is complementary, and cannot replace, the more traditional analyses of model fit. Figure 5 shows that, especially when the high-frequency component is defined as having frequency less than one year, the volatility of M_{t+1} at various frequencies is higher than, or very close to, its lower bound. Thus our model does not just demonstrate a better fit with the data, but also outperforms the CCAPM in the frequency domain. There are two more points seen in this figure that are worthy of discussion. First, for none of the alternatives is the lowest bound (close to $\omega = 1$) crossed by the SDF. This is a direct result of the fact that the filters do not capture components of the data with cycle length less than 2 months. Second, the larger the cycle length included in the high-frequency component over one year, the lower the volatility of M_{t+1} at the higher frequencies. This is because the low-frequency components begin dominating the series, and thus obscure the movements at the higher frequencies, that are lower in magnitude.[9]

We have therefore seen that considering the adjustment costs of consumption goes a good deal towards explaining the observed equity premium, in both the time and the frequency domain. To check if these results hold for other test portfolios, we check its performance with the 25 Fama-French portfolios sorted by size and book to market, before breifly confirming that the high-frequency component we have identified is not proxied by other consumption data that has been used in some studies to obtain superior fit to data.

4.1 Robustness check

Table 4 contains the results of the two-stage GMM estimation of the parameters of M_{t+1} and M_{t+1}^{CCAPM} when the 25 Fama-French portfolios sorted by size and the book-to-market ratio are the test assets, and unity and the lagged CRSP value-weighted equity return are the instruments. Given the relatively better performance of the model using consumption with cycle length less than a year, we restrict our analysis to high-frequency consumption with cycle length less than six months and one year. The estimate of risk aversion of the CCAPM is much lower, the estimate of the risk free rate and the model error are an order of magnitude higher, and the volatility of the SDF is much lower than when the equity premium is the test asset. Comparing the results in the different panels, we see that the model error is marginally lower than for the CCAPM when consumption with cycle length less than one year is defined as the high-frequency component. When consumption with cycle length less than six months is considered, the model error is marginally higher irrespective of the filter used. Based on this evidence, it is hard to distinguish between the two models. The expected risk-free rate is much higher than when the equity premium is used as the test asset, and there is little difference between the models along this measure either. M_{t+1} has a much higher volatility than M_{t+1}^{CCAPM}, irrespective of the filter or cycle length cutoff used to determine high-frequency consumption. The parameter estimates are consistent across the various methods used to determine high-frequency consumption. Taken together, there is no significant improvement achieved by distinguishing between consumption at different frequencies for these test assets.

Analysis in the frequency domain, however, does distinguish between the two models. The figure in the first row of Figure 6 shows that the CCAPM does not produce an SDF that is volatile enough to be larger than the lower bound of volatility at nearly any frequency. The four graphs in the second and third row - plotted using high-frequency data determined for cutoffs of different cycle lengths, using the aymmetric CF filter - show that the volatility of the SDF produced by considering heterogeneity of adjustment costs is sufficiently volatile to make it cross the hurdle imposed by the lower bound at every frequency. This is not entirely surprising, given the higher volatility of the SDF makes it likely that it is likely to be higher than the lower bound at some frequencies at the very least. However, as we have seen, this additional volatility turns out to be unhelpful in pricing the test assets, as we have

[9]This can also be seen in by rewriting M_{t+1} as

$$M_{t+1} = \beta \left(\frac{c^l_{t+1}}{c^l_t}\right)^{-\gamma} \left[\frac{1+\left(\frac{(c_{t+1}-c^l_{t+1})}{c^l_{t+1}}\right)^{\frac{\epsilon-1}{\epsilon}}}{1+\left(\frac{(c_t-c^l_t)}{c^l_t}\right)^{\frac{\epsilon-1}{\epsilon}}}\right]^{\frac{1-\epsilon\gamma}{\epsilon-1}}$$

At the extreme, including all consumption at all cycles in the high-frequency consumption ($c_t = c^l_t$), this reduces to

$$M_{t+1} = \beta \left(\frac{c_{t+1}}{c_t}\right)^{\gamma}$$

which is the CCAPM, and which we know has poor high-frequency performance.

seen previously.

Thus, incorporating consumption at different frequencies in the utility function helps explain the observed equity premium, but does not help - and neither does its addition do any do harm - in explaining the cross-section of equities, here represented by the 25 Fama-French portfolios. At least on this evidence, it appears that while risks of the change in the ratio of consumption of the two goods are priced in the broad market, they are not in the cross-section of stocks. Why this should be so is an interesting question that merits further study.

Finally, we check that neither of our consumption components is simply the growth rate of two other consumption goods that are known to help achieve a better model fit. Aït-Sahalia, Parker, and Yogo (2004) and Yogo (2006) show that the growth rate of luxury consumption and that of durable goods, respectively, help explain the expected stock returns. Using the data on Motohiro Yogo's website on the growth rate of luxury goods, and that of durable goods from the website of the Bureau of Economic Analysis, in Table 5 we report the growth rate of consumption of high and low adjustment cost goods defined for cycle length less than one year with that of the other variables. The growth rate of luxury goods is only available quarterly and annually, and we construct quarterly and annual growth rates of the components of consumption by aggregating monthly consumption over the required horizon and then calculating growth rates over the previous period. Panels A, B and C show that growth of the high-frequency component has little correlation with growth of the other types of consumption, irrespective of the filter used; the only exception is a marginal negative correlation between the annual growth rate of luxury sales and the high-frequency component obtained from the symmetric BK filter. This correlation is not significantly large enough in magnitude, however, to change the conclusion that the two kinds of consumption are different. Panels D, E and F repeat this exercise with the low-frequency component. The correlation with the growth of durable goods is not significant; however, there is significant correlation with the growth rate of luxury goods at the quarterly and annual horizon. Thus, it appears that luxury goods consumption shares some features of high adjustment cost goods; this appears reasonable.

All the evidence therefore points to the conclusion that the low and high adjustment cost components of consumption are not simply proxies for other well-known consumption.

5 Conclusion

Our paper is an attempt at taking a closer look inside the nondurables and services consumption that is frequently aggregated when studying its impact on the equity premium. All goods inside the consumption basket are not perfectly substitutable, and therefore, their composition is a relevant object of study. These goods are heterogenous along many dimensions; we choose one - the cost of adjustment - and show that this plays an important role in pricing the equity premium. Using the property that frequency of consumption adjustment is inversely related to costs of adjustment, we proxy the goods with low adjustment costs by the high-frequency component of aggregate consumption obtained from band pass filters. We show that splitting consumption along these lines resolves the equity premium puzzle, and also improves the frequency-domain performance with respect to the CCAPM.

An obvious extension of our analysis is a generalisation to a model incorporating goods with more than just two levels of adjustment costs. It could also be interesting to think of other possible sources of heterogeneity that might matter in asset pricing, which would be a step on the way to developing a general consumption-based asset pricing model that recognises the differences between goods that compose the consumption basket. Our message is simple: opening the consumption box may yield useful insights into asset pricing dynamics.

References

Aït-Sahalia, Y., J.A. Parker, and M. Yogo. 2004. "Luxury goods and the equity premium." *The Journal of Finance* 59 (6):2959–3004.

Bansal, R. and A. Yaron. 2004. "Risks for the long run: a potential resolution of asset pricing puzzles." *The Journal of Finance* 59 (4):1481–1509.

Baxter, M. and R.G. King. 1999. "Measuring business cycles: approximate band-pass filters for economic time series." *Review of Economics and Statistics* 81 (4):575–593.

Berkowitz, J. 2001. "Generalized spectral estimation of the consumption-based asset pricing model." *Journal of Econometrics* 104 (2):269–288.

Brockwell, P.J. and R.A. Davis. 1987. *Time series: theory and methods*. Springer Verlag.

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 (2):205–251.

Christiano, L.J. and T.J. Fitzgerald. 2003. "The band pass filter." *International Economic Review* 44 (2):435–465.

Cogley, T. 2001. "A frequency decomposition of approximation errors in stochastic discount factor models." *International Economic Review* 42 (2):473–503.

Epstein, L.G. and S.E. Zin. 1991. "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis." *Journal of Political Economy* 99 (2):263–286.

Hansen, L.P. and R. Jagannathan. 1991. "Implications of security market data for models of dynamic economies." *Journal of Political Economy* 99 (2):225–62.

Hodrick, R.J. and E.C. Prescott. 1997. "Postwar US business cycles: an empirical investigation." *Journal of Money, Credit, and Banking* 29 (1):1–16.

Marshall, D.A. and N.G. Parekh. 1999. "Can costs of consumption adjustment explain asset pricing puzzles?" *The Journal of Finance* 54 (2):623–654.

Otrok, C., B. Ravikumar, and C.H. Whiteman. 2007. "A generalized volatility bound for dynamic economies." *Journal of Monetary Economics* 54 (8):2269–2290.

Parker, J.A. and C. Julliard. 2005. "Consumption risk and the cross section of expected returns." *Journal of Political Economy* 113 (1):185–222.

Piazzesi, M., M. Schneider, and S. Tuzel. 2007. "Housing, consumption and asset pricing." *Journal of Financial Economics* 83 (3):531–569.

Rietz, T.A. 1988. "The equity risk premium: a solution." *Journal of Monetary Economics* 22 (1):117–131.

Savov, A. 2011. "Asset pricing with garbage." *The Journal of Finance* 66 (1):177–201.

Yogo, M. 2006. "A consumption-based explanation of expected stock returns." *The Journal of Finance* 61 (2):539–580.

Appendix

In this section, we show the validity of the expressions (3), (4) and (5). Set $a = (c^h_{t+1}/c^l_{t+1})^{(1/\epsilon)}$ and $b = (c^h_t/c^l_t)^{(1/\epsilon)}$. Since the first two terms of M_{t+1} and M^{CCAPM}_{t+1} are identical, we compare the third term. Thus, $M_{t+1} > M^{CCAPM}_{t+1}$ is equivalent to the expression

$$\left(\frac{1+a^{\epsilon-1}}{1+b^{\epsilon-1}}\right)^{\frac{1-\epsilon\gamma}{\epsilon-1}} > \left(\frac{1+a^{\epsilon}}{1+b^{\epsilon}}\right)^{-\gamma} \tag{6}$$

From our empirical section, we know that $a, b > 1$; further, that the two goods are substitutes ($\epsilon > 1$) and that the intertemporal elasticity of substitution (γ) is greater than one. Under these conditions, we can rewrite the condition (6) as

$$\begin{aligned} \left(\frac{1+a^{\epsilon-1}}{1+b^{\epsilon-1}}\right)^{(1-\epsilon\gamma)} &> \left(\frac{1+a^{\epsilon}}{1+b^{\epsilon}}\right)^{\gamma-\epsilon\gamma} \\ \Rightarrow \left(\frac{1+a^{\epsilon-1}}{1+b^{\epsilon-1}}\right)\left[\left(\frac{(1+a^{\epsilon})/(1+b^{\epsilon})}{(1+a^{\epsilon-1})/(1+b^{\epsilon-1})}\right)^{\epsilon}\right]^{\gamma} &> \left(\frac{1+a^{\epsilon}}{1+b^{\epsilon}}\right)^{\gamma} \end{aligned}$$

Since γ is positive, if it can be shown that

$$\left(\frac{(1+a^{\epsilon})/(1+b^{\epsilon})}{(1+a^{\epsilon-1})/(1+b^{\epsilon-1})}\right)^{\epsilon} > \left(\frac{1+a^{\epsilon}}{1+b^{\epsilon}}\right) \tag{7}$$

when $a > b$, then this is the condition for which condition (3) is satisfied, since

$$\left(\frac{1+a^{\epsilon-1}}{1+b^{\epsilon-1}}\right) > 1$$

when $a > b$. Reversing the signs in this argument allows us to conclude that $a < b$ is the condition under which (4) is satisfied, provided that we are able to prove that under this condition,

$$\left(\frac{1+a^{\epsilon}}{1+b^{\epsilon}}\right)^{\epsilon-1} < \left(\frac{1+a^{\epsilon-1}}{1+b^{\epsilon-1}}\right)^{\epsilon} \tag{8}$$

which is a restatement of (7) with a change in the direction of the inequality.

While we are unable to prove this for all values of ϵ, γ, we provide evidence that this is true for reasonable values of these parameters based on our empirical analysis; and of the a and b, as observed in the data. Using the results from the empirical section, we find that $a, b \in [9.7, 27.4]$ and $k = (a/b) \in [0.78, 1.44]$; further, that $\epsilon > 1$. We therefore check that the condition $a/b > 1$ ($a/b < 1$) determines that M_{t+1}/M^{CCAPM}_{t+1} is greater (less) than 1 for $b \in [8, 50]$ and various values of ϵ. We do this by choosing a fixed value of $\epsilon = 1.5$ and plotting the ratio of the two sides of expression (8) (left hand side / right hand side) for the aforementioned values of b. We plot separate graphs for $k \in [0.5, 1]$ and $k \in [1, 2]$ since plotting them on the same graphs makes it difficult to see that this difference is less than one for $k \in [0.5, 1]$. The reason is that the values encountered in this range are much smaller in magnitude than the values greater than one encountered in the range $k \in [1, 2]$. We repeat these plots for $\epsilon = 3$ and $\epsilon = 5$ to confirm that the pattern remains the same. All these plots, collected in Figure 1, confirm the conditions shown in (3) and (4). (5) is easily seen by substituting $a = b$ in equations (1) and (2).

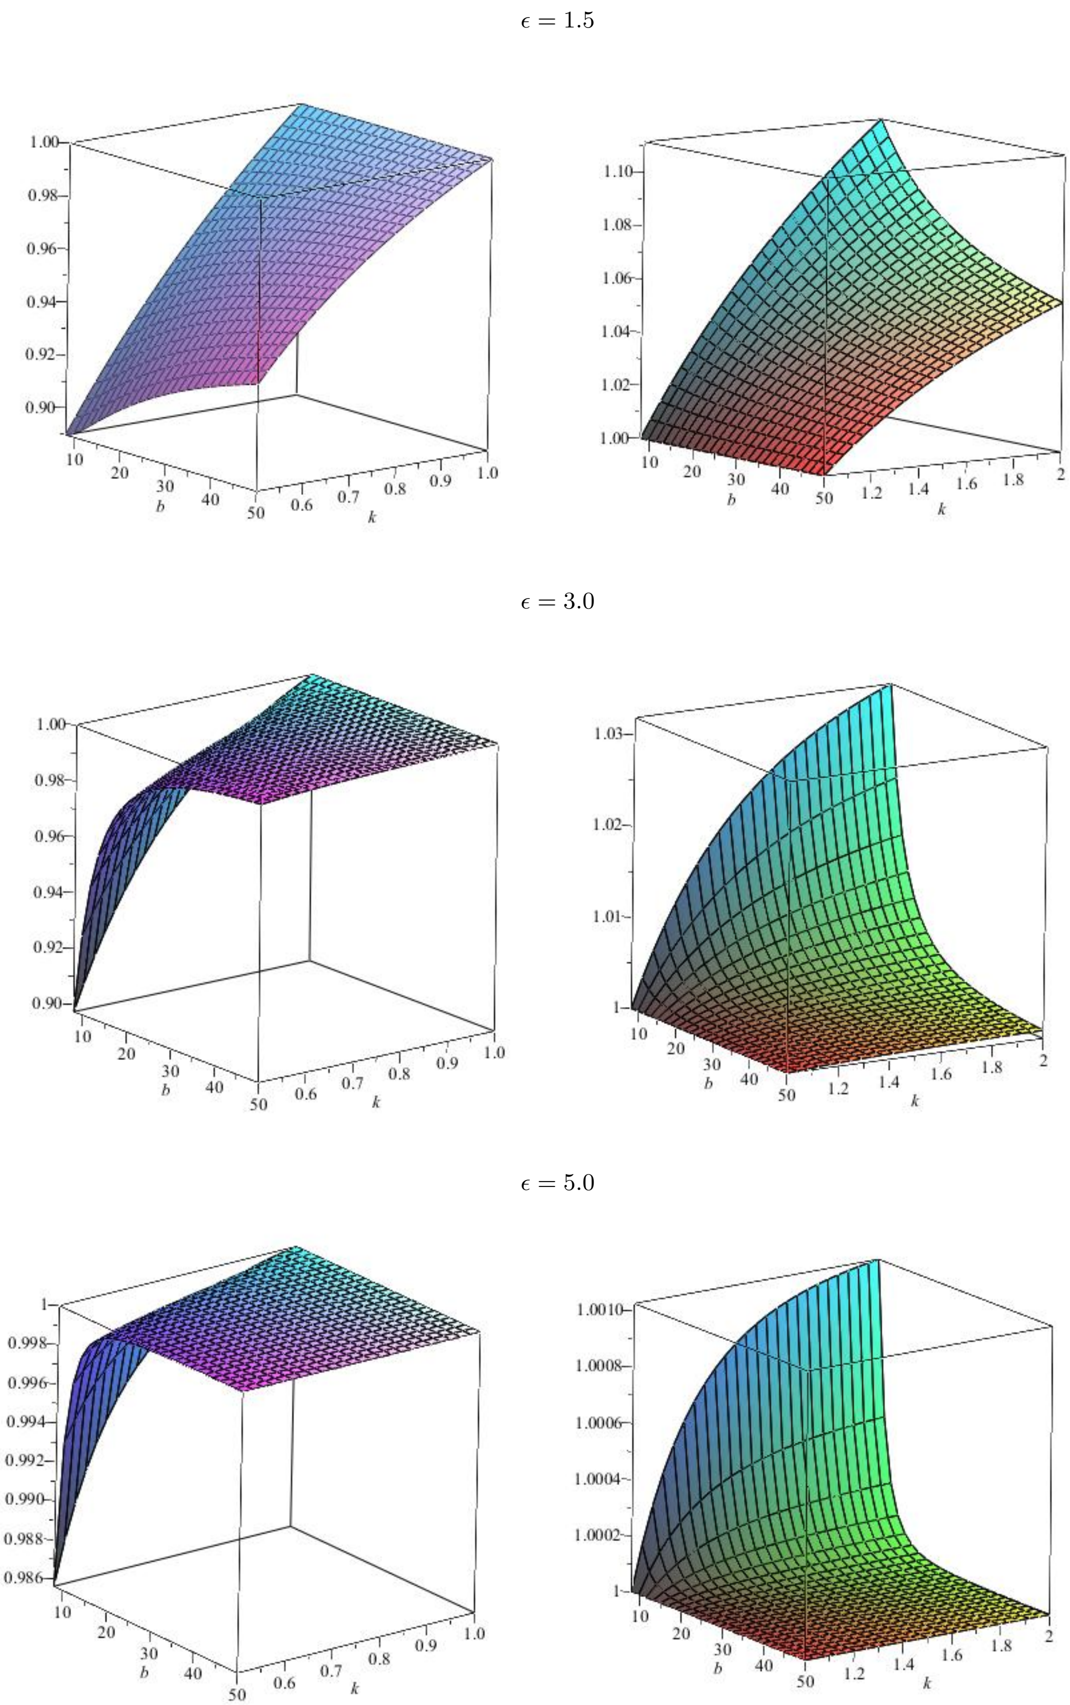

Figure 1: Comparison of SDF of CCAPM and that of our modification
This set graph plots a measure of the ratio between the SDFs of our modification and of the CCAPM for plausible values of c_t^h/c_t^l (b) and the ratio $(c_{t+1}^h/c_{t+1}^l)/(c_t^h/c_t^l)$ (k). Each row uses a different value of ϵ; the left-hand side columns plot the graphs for $k \in [0.5, 1.0]$ and the right-hand plots for $k \in [1.0, 2.0]$.

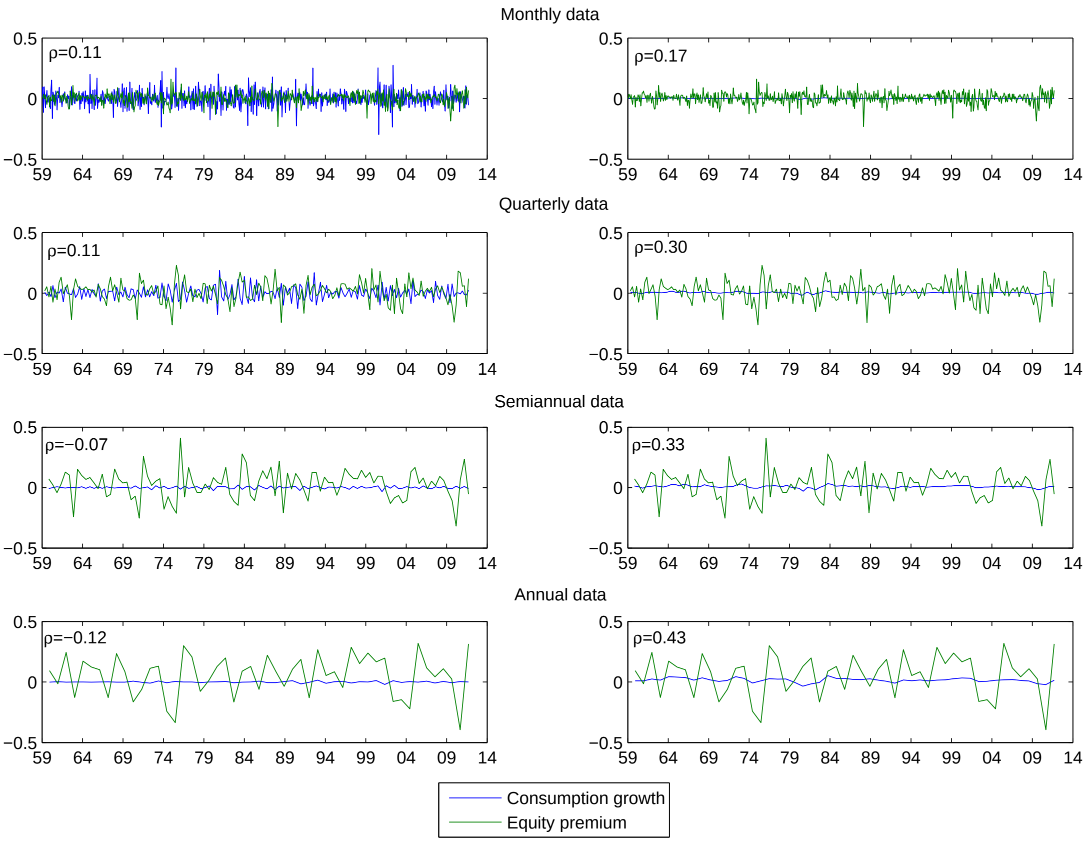

Figure 2: Growth of high- and low-frequency components of consumption and of the equity premium
The four left-(right-)hand plots graph the time series of growth of the high-(low-)frequency component of consumption and the equity premium for the entire sample. Each row has a different horizon for which the growth is calculated, indicated in the heading. The correlation between the two time series is indicated in the top-left corner of each graph. The high-frequency component is obtained using the asymmetric filter in Christiano and Fitzgerald (2003) on monthly nondurable and services consumption data from January 1959 to December 2010; the residual is the low-frequency component. The consumption series are aggregated over the period specified and growth rates are calculated over the preceding period.

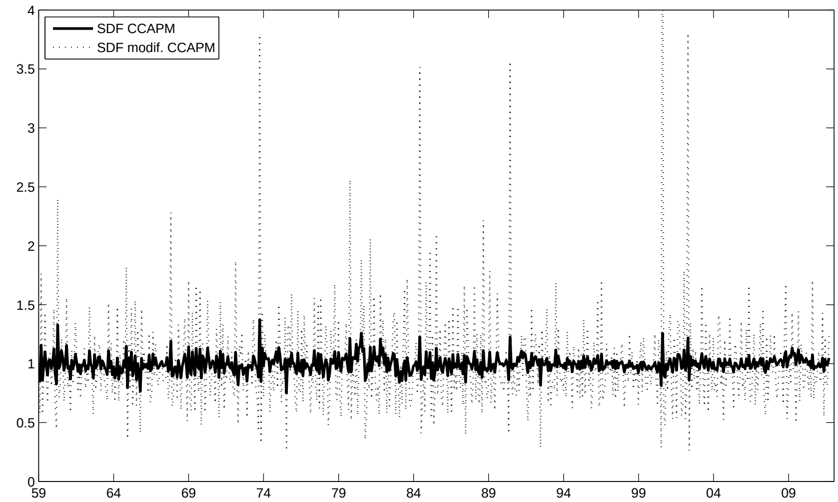

Figure 3: Time series of SDF for CCAPM and modified CCAPM using the asymmetric CF filter
This graph plots the estimated realisations of the stochastic discount factor at each month beginning February 1959 up to December 2010. The SDF is estimated for the standard consumption CAPM and the modified version we propose in the text, where consumption is filtered as using the asymmetric filter proposed in Christiano and Fitzgerald (2003). The component of consumption that has frequency less than one year is labeled 'high-frequency'. The test assets used in model estimation are the equity premium and the yield on the one month T-Bill; the instruments used are unity and the lagged equity premium.

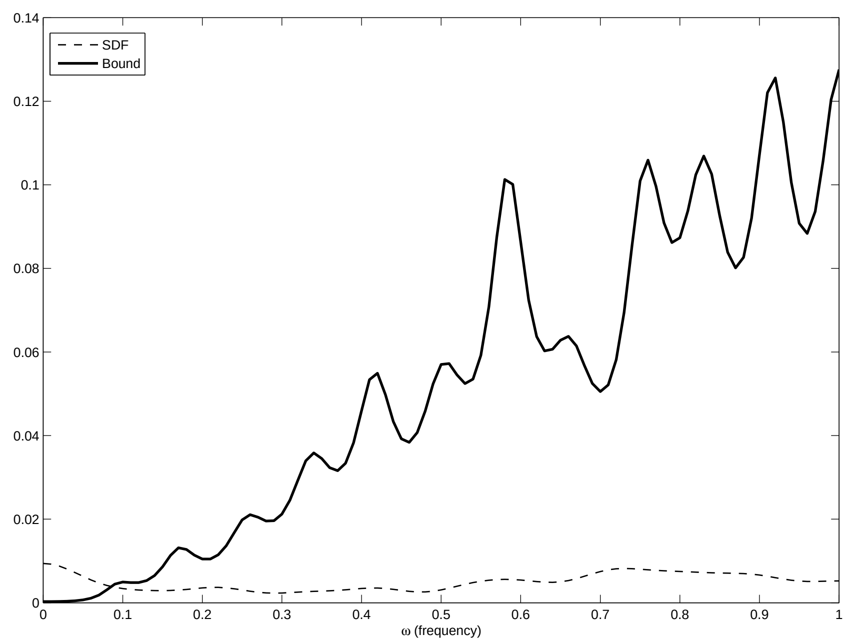

Figure 4: Spectrum of SDF and its lower bound for CCAPM
This graph plots the spectrum of the stochastic discount factor obtained from the estimation of the CCAPM and its lower bound, using the procedure detailed in Otrok, Ravikumar, and Whiteman (2007). The model is estimated on data at monthly frequency for the period January 1959 - December 2010. The test assets used in model estimation are the equity premium and the yield on the one month T-Bill; the instruments used are unity and the lagged equity premium. The spectrum is estimated for 101 values of frequency and is smoothed using Bartlett weights on 24 leads and lags.

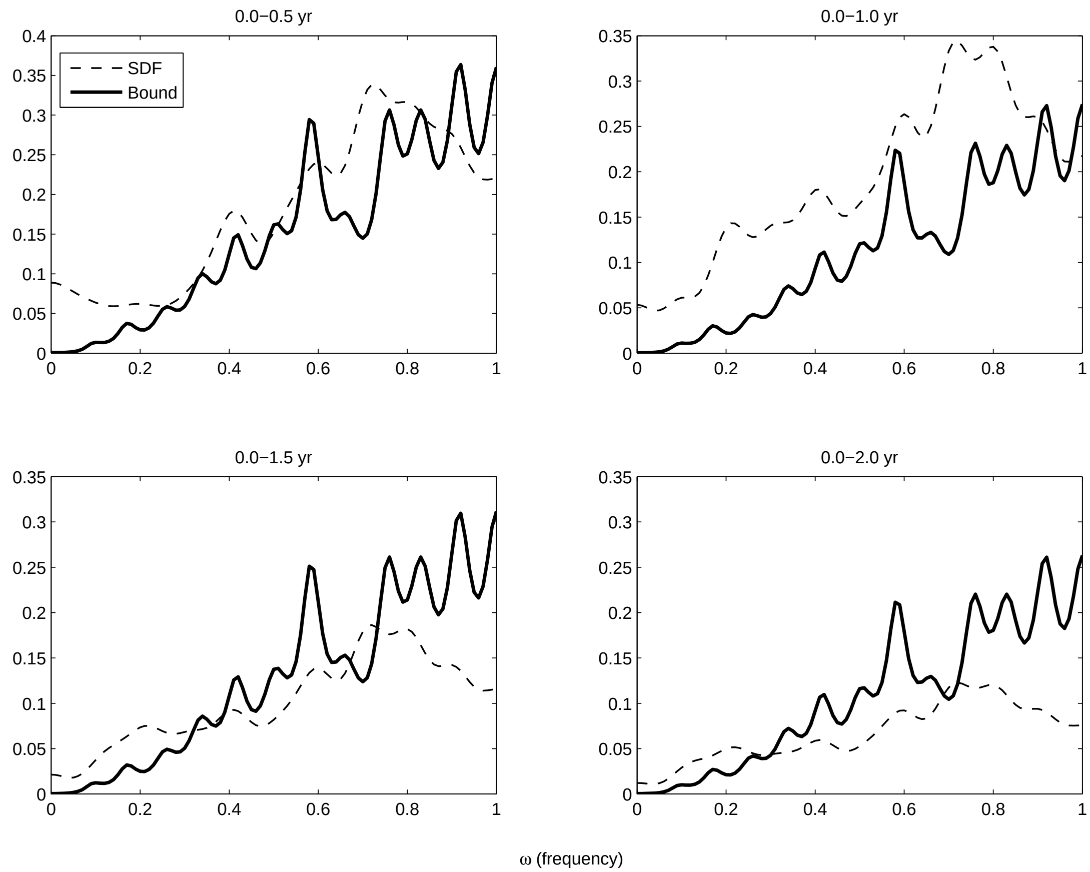

Figure 5: Spectrum of SDF and its lower bound for modified CCAPM using the asymmetric CF filter
This graph plots the spectrum of the stochastic discount factor obtained from the estimation of the modified CCAPM we propose in the text where the high- and low-frequency consumption is obtained by using the asymmetric filter proposed in Christiano and Fitzgerald (2003) on real monthly consumption data for the period January 1959 - December 2010, using the procedure detailed in Otrok, Ravikumar, and Whiteman (2007). The title of each panel indicates the frequency which is taken to define the high-frequency component of consumption. The test assets used in model estimation are the equity premium and the yield on the one month T-Bill; the instruments used are unity and the lagged equity premium. The spectrum is estimated for 101 discrete values of ω and is smoothed using Bartlett weights on 24 leads and lags.

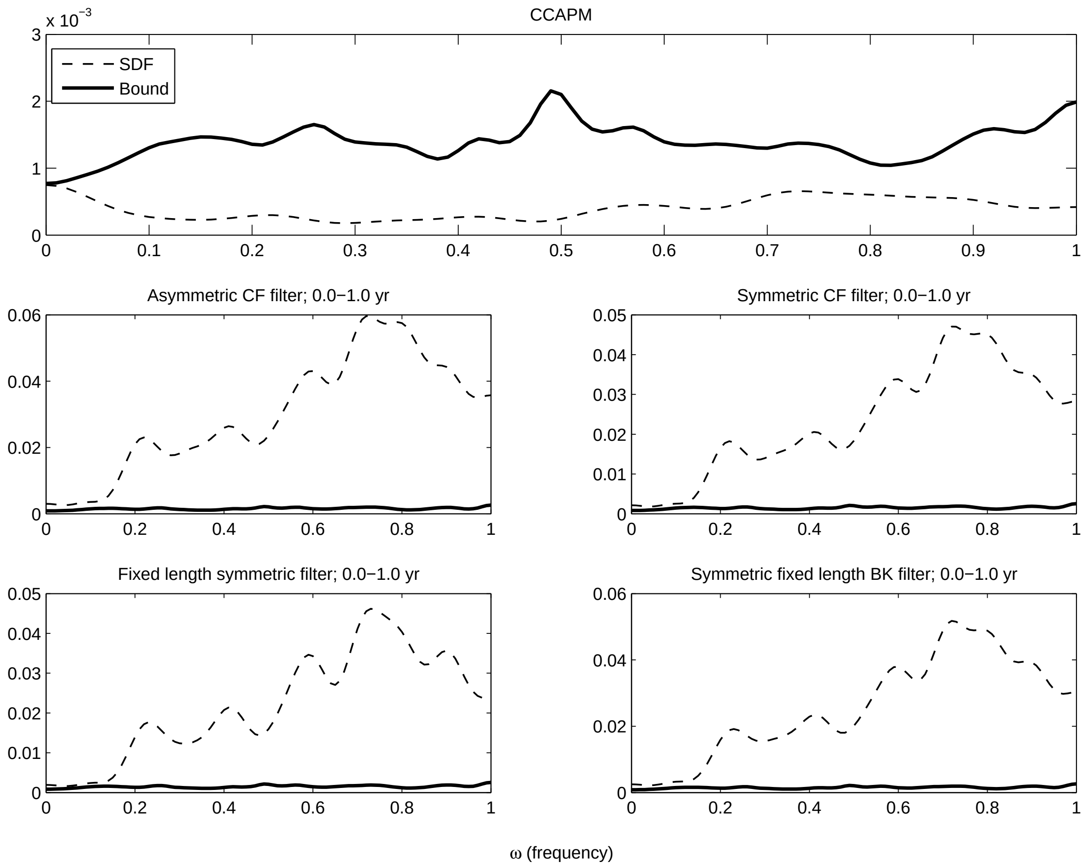

Figure 6: Spectrum of SDF and its lower bound for the 25 Fama-French portfolios
This graph plots the spectrum of the stochastic discount factor obtained from the estimation of the CCAPM and the modified CCAPM we propose where the high-frequency consumption is obtained by filtering real monthly consumption data for the period January 1959 - December 2010 of the frequency less than one year, using the procedure detailed in Otrok, Ravikumar, and Whiteman (2007). The title of each panel indicates the type of filter used. The test assets used in model estimation are 25 Fama-French portfolios; the instruments used are unity and the lagged equity premium. The spectrum is estimated for 101 values of frequency and is smoothed using Bartlett weights on 24 leads and lags.

	High-frequency consumption				Low frequency consumption			
	0-0.5y	0-1.0y	0-1.5y	0-2.0y	0.5y+	1.0y+	1.5y+	2.0y+
Panel A: Asymmetric Christiano-Fitzgerald filter								
μ	125.00	124.99	125.01	124.98	2190.42	2190.42	2190.40	2190.44
σ	4.86	6.29	7.35	8.17	482.51	482.50	482.49	482.51
ρ_1	-0.38	0.07	0.30	0.42	1.00	1.00	1.00	1.00
ρ_2	-0.27	-0.11	0.09	0.23	0.99	0.99	0.99	0.99
ρ_3	0.09	-0.15	-0.03	0.09	0.99	0.99	0.99	0.99
ρ_6	0.03	-0.12	-0.30	-0.28	0.97	0.97	0.97	0.97
ρ_{12}	-0.09	-0.07	0.02	-0.15	0.94	0.94	0.94	0.94
μ^g	0.20%	0.23%	0.24%	0.24%	0.13%	0.13%	0.13%	0.13%
σ^g	6.38%	6.77%	6.95%	7.06%	0.23%	0.18%	0.17%	0.16%
ρ^g_1	-0.54	-0.40	-0.35	-0.33	0.88	0.97	0.99	0.99
ρ^g_2	-0.09	-0.08	-0.05	-0.04	0.59	0.90	0.96	0.97
ρ^g_3	0.17	0.08	0.09	0.10	0.30	0.79	0.90	0.93
ρ^g_6	-0.01	-0.02	-0.04	-0.04	0.20	0.40	0.66	0.75
ρ^g_{12}	-0.09	-0.09	-0.08	-0.09	0.12	0.20	0.17	0.27
Panel B: Symmetric Christiano-Fitzgerald filter								
μ	125.00	124.99	125.01	124.95	2190.42	2190.42	2190.40	2190.47
σ	4.85	6.29	7.35	8.19	482.51	482.50	482.48	482.49
ρ_1	-0.39	0.07	0.30	0.43	1.00	1.00	1.00	1.00
ρ_2	-0.27	-0.11	0.09	0.24	0.99	0.99	0.99	0.99
ρ_3	0.10	-0.15	-0.03	0.10	0.99	0.99	0.99	0.99
ρ_6	0.02	-0.12	-0.30	-0.28	0.97	0.97	0.97	0.97
ρ_{12}	-0.09	-0.07	0.02	-0.15	0.94	0.94	0.94	0.94
μ^g	0.20%	0.23%	0.24%	0.25%	0.13%	0.13%	0.13%	0.13%
σ^g	6.38%	6.78%	6.94%	7.06%	0.23%	0.18%	0.17%	0.16%
ρ^g_1	-0.54	-0.40	-0.35	-0.33	0.88	0.96	0.96	0.96
ρ^g_2	-0.09	-0.07	-0.05	-0.03	0.60	0.89	0.93	0.94
ρ^g_3	0.17	0.08	0.09	0.09	0.30	0.78	0.89	0.91
ρ^g_6	-0.01	-0.02	-0.04	-0.04	0.20	0.39	0.66	0.73
ρ^g_{12}	-0.10	-0.09	-0.08	-0.09	0.12	0.20	0.17	0.27
Panel C: Fixed-length symmetric filter								
μ	125.00	125.00	124.99	124.96	2190.42	2190.42	2190.42	2190.45
σ	4.86	6.27	7.13	7.80	482.51	482.50	482.58	482.46
ρ_1	-0.39	0.05	0.26	0.38	1.00	1.00	1.00	1.00
ρ_2	-0.28	-0.13	0.06	0.19	0.99	0.99	0.99	0.99
ρ_3	0.11	-0.15	-0.07	0.04	0.99	0.99	0.99	0.99
ρ_6	0.01	-0.09	-0.26	-0.29	0.97	0.97	0.97	0.97
ρ_{12}	-0.18	-0.20	0.09	0.09	0.94	0.94	0.94	0.94
μ^g	0.20%	0.23%	0.23%	0.24%	0.13%	0.13%	0.13%	0.13%
σ^g	6.33%	6.77%	6.82%	6.87%	0.24%	0.20%	0.19%	0.18%
ρ^g_1	-0.54	-0.40	-0.35	-0.33	0.71	0.69	0.66	0.56
ρ^g_2	-0.09	-0.09	-0.05	-0.03	0.52	0.70	0.71	0.64
ρ^g_3	0.17	0.08	0.07	0.08	0.33	0.71	0.77	0.73
ρ^g_6	-0.02	-0.02	-0.05	-0.05	0.18	0.39	0.55	0.56
ρ^g_{12}	-0.15	-0.18	0.01	0.08	0.08	0.15	0.15	0.20
Panel D: Fixed-length symmetric Baxter-King filter								
μ	125.00	125.01	124.98	124.93	2190.41	2190.40	2190.43	2190.48
σ	4.88	6.36	7.12	8.46	482.56	482.60	482.50	482.29
ρ_1	-0.37	0.09	0.26	0.46	1.00	1.00	1.00	1.00
ρ_2	-0.26	-0.08	0.05	0.29	0.99	0.99	0.99	0.99
ρ_3	0.12	-0.10	-0.07	0.17	0.99	0.99	0.99	0.99
ρ_6	0.02	-0.06	-0.26	-0.14	0.97	0.97	0.97	0.97
ρ_{12}	-0.14	-0.11	0.02	-0.08	0.94	0.94	0.94	0.94
μ^g	0.20%	0.23%	0.23%	0.24%	0.13%	0.13%	0.13%	0.13%
σ^g	6.31%	6.70%	6.85%	6.94%	0.24%	0.20%	0.18%	0.16%
ρ^g_1	-0.54	-0.40	-0.35	-0.32	0.72	0.73	0.66	0.59
ρ^g_2	-0.09	-0.08	-0.05	-0.04	0.53	0.73	0.71	0.67
ρ^g_3	0.17	0.08	0.08	0.09	0.33	0.74	0.78	0.79
ρ^g_6	-0.02	-0.01	-0.05	-0.05	0.19	0.41	0.57	0.64
ρ^g_{12}	-0.11	-0.10	-0.05	-0.07	0.06	0.11	0.19	0.31

Table 1: Summary statistics of filtered consumption data

This table presents time series and cross-sectional statistics of the filtered real monthly nondurables and services consumption from 1959 to 2010. Each panel presents results obtained using a different filtering method. Columns 2-5 (6-9) show statistics for high-(low-)frequency component of consumption, with each column using a different definition of high-(low-) frequencies. For Panels C and D, the length is fixed at 12 months. Within each panel, the subscript g indicates that the statistics are for the growth of filtered consumption series, with μ, σ, ρ_i representing the mean, standard deviation and autocorrelation of lag i, respectively.

	Asy CF; 0.5yr	Sym CF; 0.5yr	Fix len; 0.5yr	Sym BK; 0.5yr	Asy CF; 1yr	Sym CF; 1yr	Fix len; 1yr
Symmetric CF filter; 0.5yr	0.9995						
Fixed length filter; 0.5yr	0.9916	0.9924					
Symmetric BK filter; 0.5yr	0.9844	0.9851	0.9924				
Asymmetric CF filter; 1yr	0.7731	0.7734	0.7744	0.7666			
Symmetric CF filter; 1yr	0.7728	0.7731	0.7741	0.7662	0.9996		
Fixed length filter; 1yr	0.7819	0.7825	0.7857	0.7791	0.9851	0.9848	
Symmetric BK filter; 1yr	0.7651	0.7655	0.7682	0.786	0.9636	0.9633	0.9804

Table 2: Correlation between high-frequency consumption obtained using alternate filters
This table presents the correlation between the high-frequency consumption obtained by the four filters discussed in the text by filtering out all frequencies above six months, and one year. The data is of monthly frequency from January 1959 to December 2010.

	β	γ	ϵ	δ	R_{t+1} (%)	σ_{SDF}
All freq	1.0150	16.7417		0.0692	0.27	0.0723
	1.26E-05	7.76E-08		(4.29E-10)		
Panel A: Asymmetric CF filter						
0.0-0.5 yr	0.9757	43.4524	3.3599	0.0617	0.16	0.4374
	1.3E-05	1.9E-10	4.9E-09	(5.63E-10)		
0.0-1.0 yr	0.9325	5.2388	0.1724	0.0595	0.19	0.4376
	4.5E-05	5.4E-08	2.2E-18	(1.15E-09)		
0.0-1.5 yr	0.9576	4.0543	0.2267	0.0586	0.18	0.3198
	1.4E-06	5.8E-08	1.9E-17	(1.53E-09)		
0.0-2.0 yr	0.9688	3.3992	0.2639	0.0621	0.20	0.2603
	2.2E-05	4.5E-09	1.3E-16	(4.94E-10)		
Panel B: Symmetric CF filter						
0.0-0.5 yr	0.9768	44.4103	3.4609	0.0618	0.16	0.4375
	1.2E-05	1.7E-10	4.9E-09	(5.53E-10)		
0.0-1.0 yr	0.9437	9.0319	1.0056	0.0604	0.18	0.4087
	3.0E-06	2.7E-12	7.0E-10	(8.69E-10)		
0.0-1.5 yr	0.9587	3.9999	0.2301	0.0587	0.18	0.3147
	6.1E-06	2.5E-08	5.7E-17	(1.45E-09)		
0.0-2.0 yr	0.9685	3.4290	0.2618	0.0622	0.20	0.2626
	1.9E-05	4.0E-09	1.2E-16	(4.78E-10)		
Panel C: Fixed length symmetric filter						
0.0-0.5 yr	0.9949	41.4096	4.0257	0.0645	0.17	0.3544
	9.3E-04	1.9E-10	1.1E-10	(2.30E-10)		
0.0-1.0 yr	0.9650	17.1998	1.6979	0.0620	0.18	0.3540
	3.9E-05	1.9E-11	1.4E-10	(5.15E-10)		
0.0-1.5 yr	0.9618	10.0064	1.2199	0.0610	0.18	0.3334
	1.7E-06	1.3E-12	6.0E-10	(7.13E-10)		
0.0-2.0 yr	0.9589	4.0454	0.2257	0.0618	0.16	0.3153
	1.3E-05	1.3E-09	1.2E-17	(5.53E-10)		
Panel D: Symmetric fixed length BK filter						
0.0-0.5 yr	0.9959	42.4755	4.1642	0.0644	0.18	0.3558
	8.9E-04	1.8E-10	1.1E-10	(2.36E-10)		
0.0-1.0 yr	0.9700	23.0372	2.1016	0.0619	0.19	0.3626
	2.0E-05	7.8E-10	8.9E-09	(5.38E-10)		
0.0-1.5 yr	0.9563	4.1810	0.2186	0.0610	0.18	0.3269
	6.9E-04	1.4E-09	1.7E-17	(7.10E-10)		
0.0-2.0 yr	0.9756	3.0462	0.2905	0.0635	0.19	0.2260
	7.4E-04	9.4E-09	4.9E-15	(3.19E-10)		

Table 3: Model comparison of equity premium

This table presents the estimates of the parameters $(\beta, \gamma, \epsilon)$, the goodness of fit (δ), the risk-free rate (R_f), and the standard deviation of the stochastic discount factor (σ_{SDF}) of the CCAPM and our proposed alternative when the equity premium and the one month risk-free rate are the test assets. The standard errors of the parameter estimates are shown directly below the estimate. Values in parantheses are the p values of the χ^2 test that the model error is zero. The first two rows show the parameter estimates of the CCAPM. Panels A-D show the estimates of the version of the CCAPM we propose that takes into account consumption adjustment costs, with each panel containing estimates of the high and low frequencies of consumption obtained through a different filter; CF refers to the filter proposed by Christiano and Fitzgerald (2003) and BK refers to that suggested by Baxter and King (1999). Each row within a panel uses a different definition of the high-frequency component of consumption. All filters are used over monthly real nondurables and services consumption data from January 1959 to December 2010. The estimates are obtained using a two-stage GMM with the two instrumental variables: unity, and the lagged equity returns.

	β	γ	ϵ	δ	$R_{f,t+1}$ (%)	σ_{SDF}
All freq	0.994	4.735		0.603	1.19	0.020
	4.84E-02	3.32E-07		(2.28E-52)		
Panel A: Asymmetric CF filter						
0.0-0.5 yr	0.973	2.992	0.720	0.606	1.37	0.174
	5.33E-02	1.09E-08	2.52E-07	(4.09E-53)		
0.0-1.0 yr	0.973	2.736	0.695	0.600	1.36	0.174
	1.05E-03	4.25E-08	1.28E-08	(1.89E-52)		
Panel B: Symmetric CF filter						
0.0-0.5 yr	0.973	3.000	0.720	0.605	1.35	0.174
	5.10E-02	1.03E-08	2.35E-07	(4.39E-53)		
0.0-1.0 yr	0.977	2.371	0.680	0.599	1.21	0.154
	1.61E-03	6.75E-08	3.48E-08	(2.58E-52)		
Panel C: Fixed length symmetric filter						
0.0-0.5 yr	0.972	3.199	0.722	0.604	1.38	0.184
	4.43E-02	1.09E-08	2.01E-07	(6.68E-53)		
0.0-1.0 yr	0.978	2.282	0.681	0.599	1.21	0.148
	1.40E-03	6.97E-08	4.02E-08	(2.67E-52)		
Panel D: Symmetric fixed length BK filter						
0.0-0.5 yr	0.971	3.234	0.724	0.604	1.39	0.185
	3.90E-02	9.06E-07	1.60E-07	(6.80E-53)		
0.0-1.0 yr	0.976	2.535	0.699	0.598	1.26	0.160
	1.14E-03	5.10E-08	1.92E-08	(2.91E-52)		

Table 4: Model comparison with 25 Fama-French portfolios as test assets
This table presents the estimates of the parameters $(\beta, \gamma, \epsilon)$, the goodness of fit (δ), the risk-free rate (R_f), and the standard deviation of the stochastic discount factor (σ_{SDF}) of the CCAPM and our proposed alternative using the 25 Fama-French portfolios as test assets. The standard errors of the parameter estimates are shown directly below the estimate. The first two rows show the parameter estimates of the CCAPM. Values in parantheses are the p values of the χ^2 test that the model error is zero. Panels A-D show the estimates of the version of the CCAPM we propose that takes into account consumption adjustment costs, with each panel containing estimates of the high and low frequencies of consumption obtained through a different filter; CF refers to the filter proposed by Christiano and Fitzgerald (2003) and BK refers to that suggested by Baxter and King (1999). Each row within a panel uses a different definition of the high-frequency component of consumption. All filters are used over monthly real nondurables and services consumption data from January 1959 to December 2010. The estimates are obtained using a two-stage GMM with the two instrumental variables: unity, and the lagged equity returns.

	Asymmetric CF filter	Symmetric CF filter	Fixed length filter	Symmetric BK filter
High-frequency consumption				
Panel A: Monthly growth rates				
Symmetric CF filter	0.9999			
Fixed length filter	0.9955	0.9955		
Symmetric BK filter	0.9976	0.9975	0.9978	
Per-capita durables	-0.0260	-0.0261	-0.0325	-0.0286
Panel B: Quarterly growth rates				
Symmetric CF filter	0.9997			
Fixed length filter	0.9868	0.9864		
Symmetric BK filter	0.9846	0.9847	0.9957	
Luxury good sales	0.1270	0.1270	0.1177	0.0923
Panel C: Annual growth rates				
Symmetric CF filter	0.9990			
Fixed length filter	0.9837	0.9835		
Symmetric BK filter	0.7044	0.7055	0.8061	
Luxury sales growth	-0.0279	-0.0203	-0.0609	-0.2371
Low-frequency consumption				
Panel D: Monthly growth rates				
Symmetric CF filter	0.9946			
Fixed length filter	0.8924	0.8874		
Symmetric BK filter	0.9044	0.899	0.9913	
Per-capita durables	-0.0169	-0.0176	0.003	-0.0035
Panel E: Quarterly growth rates				
Symmetric CF filter	0.9998			
Fixed length filter	0.99	0.9898		
Symmetric BK filter	0.9887	0.9887	0.9978	
Luxury good sales	0.5617	0.5609	0.5776	0.5786
Panel E: Annual growth rates				
Symmetric CF filter	1.0000			
Fixed length filter	1.0000	1.0000		
Symmetric BK filter	0.9994	0.9994	0.9996	
Luxury sales growth	0.4286	0.4281	0.4281	0.4266

Table 5: Correlation of high-frequency consumption growth with alternate consumption growth measures
This table presents the correlation of the growth of the high-frequency consumption, obtained by filtering out the components of consumption with frequency greater than one year using the four filters described in the text, and three alternate measures of consumption that have been used to explain the cross-section of equity returns: the growth of per-capita durables consumption, and of luxury good sales. Each panel contains correlations at a different frequency, as mentioned in the heading: monthly, quarterly, and annual. Correlations for the monthly growth rates are for the period February 1959 to December 2010; for the quarterly growth rates from Q1 1987 to Q4 2001; and for annual growth rates from 1961 to 2001. Data on per-capita durables is obtained from the US Bureau of Economic Analaysis, and on luxury good sales from Motohiro Yogo's website.