



Universität St.Gallen

## C-CAPM Refinements and the Cross- Section of Returns

Paul Söderlind

January 2006 Discussion Paper 07-2006

Editor:

Prof. Jörg Baumberger  
University of St. Gallen  
Department of Economics  
Bodanstr. 1  
CH-9000 St. Gallen  
Phone +41 71 224 22 41  
Fax +41 71 224 28 85  
Email [joerg.baumberger@unisg.ch](mailto:joerg.baumberger@unisg.ch)

Publisher:

Department of Economics  
University of St. Gallen  
Bodanstrasse 8  
CH-9000 St. Gallen  
Phone +41 71 224 23 25  
Fax +41 71 224 22 98

Electronic Publication:

<http://www.vwa.unisg.ch>

# C-CAPM Refinements and the Cross-Section of Returns

Paul Söderlind

Author's address:

Professor Paul Söderlind, PhD  
Swiss Institute of Banking and Finance  
Rosenbergstr 52  
9000 St. Gallen  
Tel. +41 71 224 7080  
Fax +41 71 224 7088  
Email [Paul.Soderlind@unisg.ch](mailto:Paul.Soderlind@unisg.ch)  
Website [www.sbf.unisg.ch](http://www.sbf.unisg.ch)

## **Abstract**

This paper studies if the consumption-based asset pricing model can explain the cross-section of expected returns. The CRRA model and several refinements (habit persistence and idiosyncratic shocks) all imply that the conditional expected return is linearly increasing in the asset's conditional covariance with consumption growth. Results from quarterly data on the 25 Fama-French portfolios suggest that the model has serious problems: there are large and systematic pricing errors. In addition, the estimated time-varying effective risk aversion coefficients appear implausible and are unrelated with most candidates for habit persistence and idiosyncratic risk.

## **Keywords**

consumption-based asset pricing, habit persistence, idiosyncratic risk, conditional asset pricing

## **JEL Classification**

G12, E130, E320

# 1 Introduction

It is well documented that the standard consumption-based capital asset pricing model (C-CAPM) has serious problems with matching the high historical average excess return on equity (the “equity premium puzzle”).<sup>1</sup> This has spurred a large literature that attempts to refine the model. In particular, several papers have argued that we should consider adding idiosyncratic risk, recalibrate the risk aversion parameter, or assume habit persistence.<sup>2</sup>

This paper evaluates these suggestions by studying the cross-sectional performance of a conditional asset pricing model. The basic idea is that several refinements of C-CAMP suggest that the conditional risk premium should—in a cross-section of assets—be linearly related to the conditional covariance with consumption growth, but with a time-varying slope coefficient (the “effective” or “local” risk aversion).

The main results are that the conditional model generates pricing errors of the same order of magnitude as the unconditional model (which are large), and that the estimated slope coefficients appear to be more or less uncorrelated with macro variables, except possibly with survey measures of growth and inflation uncertainty.

The outline of the rest of the paper is as follows: Section 2 derives the relation between the expected return and covariance in the basic model and several refinements; Section 3 presents the results for the 25 Fama and French (1993) portfolios; and Section 4 concludes.

## 2 C-CAPM: Basic Model and Refinements

### 2.1 Basic Framework

This section derives a simple relation between an asset’s risk premium and its covariance with aggregate consumption growth.

The basic asset pricing equation says

$$E_{t-1}(Z_{it}M_t) = 0, \tag{1}$$

where  $Z_{it}$  is the excess return of holding asset  $i$  from period  $t - 1$  to  $t$ ,  $M_t$  is a stochastic

---

<sup>1</sup>See, for instance, Campbell (2001) and Cochrane (2001) for overviews.

<sup>2</sup>See, for instance, Mankiw (1986) and Constantinides and Duffie (1996), Epstein and Zin (1989), and Campbell and Cochrane (1999).

discount factor (SDF), and  $E_{t-1}$  denotes the expectations (conditional on the information set in  $t - 1$ ). To simplify the analysis, assume that the excess return,  $Z_{it}$ , and the log SDF,  $\ln M_t$ , have a bivariate normal distribution. Use Stein's lemma (see Appendix B) and rearrange (1) to express the risk premium (expected excess return) as

$$E_{t-1}(Z_{it}) = -\text{Cov}_{t-1}(Z_{it}, \ln M_t). \quad (2)$$

We can relax the assumption that the excess return is normally distributed: (2) holds also if  $Z_{it}$  and  $\ln M_t$  have a bivariate mixture normal (conditional) distribution—provided  $\ln M_t$  has the same mean and variance in all the mixture components (see Appendix B). This restricts the log discount factor to have a normal distribution, but allows the excess return to have a distribution with fat tails and skewness. (This is discussed in some detail below.)

The unconditional version of (2) is obtained by taking unconditional expectations of (1) and assuming that we can apply Stein's lemma (possibly extended as above) on the unconditional distribution to get

$$E(Z_{it}) = -\text{Cov}(Z_{it}, \ln M_t). \quad (3)$$

The next couple of sections discusses a number of different models of the stochastic discount factor: the standard CRRA model, Epstein-Zin recursive utility, a model with habit persistence, and a model with idiosyncratic risk.

## 2.2 The Standard CRRA Model

With constant relative risk aversion (CRRA), the SDF has the form  $M_t = (C_t/C_{t-1})^{-\gamma}$ , where  $\gamma$  is the risk aversion parameter and  $C_t$  is the consumption level.<sup>3</sup> The log SDF is therefore

$$\ln M_t = -\gamma \Delta c_t, \quad (4)$$

where  $\Delta c_t$  is the growth rate of consumption, that is,  $\Delta c_t = \ln(C_t/C_{t-1})$ . With identical investors, consumption growth equals aggregate consumption growth.

---

<sup>3</sup>A time discount rate would cancel in (1) and is therefore suppressed.

Using (4) in (3) and (2) gives

$$E(Z_{it}) = \text{Cov}(Z_{it}, \Delta c_t)\gamma, \text{ and} \quad (5)$$

$$E_{t-1}(Z_{it}) = \text{Cov}_{t-1}(Z_{it}, \Delta c_t)\gamma. \quad (6)$$

Alternatively, this could be rearranged in beta form.<sup>4</sup>

The intuition for these expressions is that an asset that has a high payoff when consumption is high, that is, when marginal utility is low, is considered risky and will require a risk premium.

A key implication of (5) is that, in a cross-section of assets, the expected return is linearly related to the covariance with aggregate consumption growth. The conditional expression, (6), is similar, except that it requires information on the conditional means and covariances. I will later demonstrate that recent refinements of the consumption-based model (habit persistence and idiosyncratic risk) also give expressions similar to the conditional expression (6), except that the effective risk aversion coefficient might be different in different time periods.<sup>5</sup>

### The Gains and Losses from Using Stein's Lemma

The gain from using (the extended) Stein's lemma is that the unknown relative risk aversion,  $\gamma$ , does not enter the covariances. This facilitates the empirical analysis considerably.<sup>6</sup>

The price of using (the extended) Stein's lemma is that we have to assume that consumption growth is normally distributed and that the excess return have a mixture normal distribution. The latter is not much of a price, since a mixture normal can take many shapes and have both skewness and excess kurtosis.

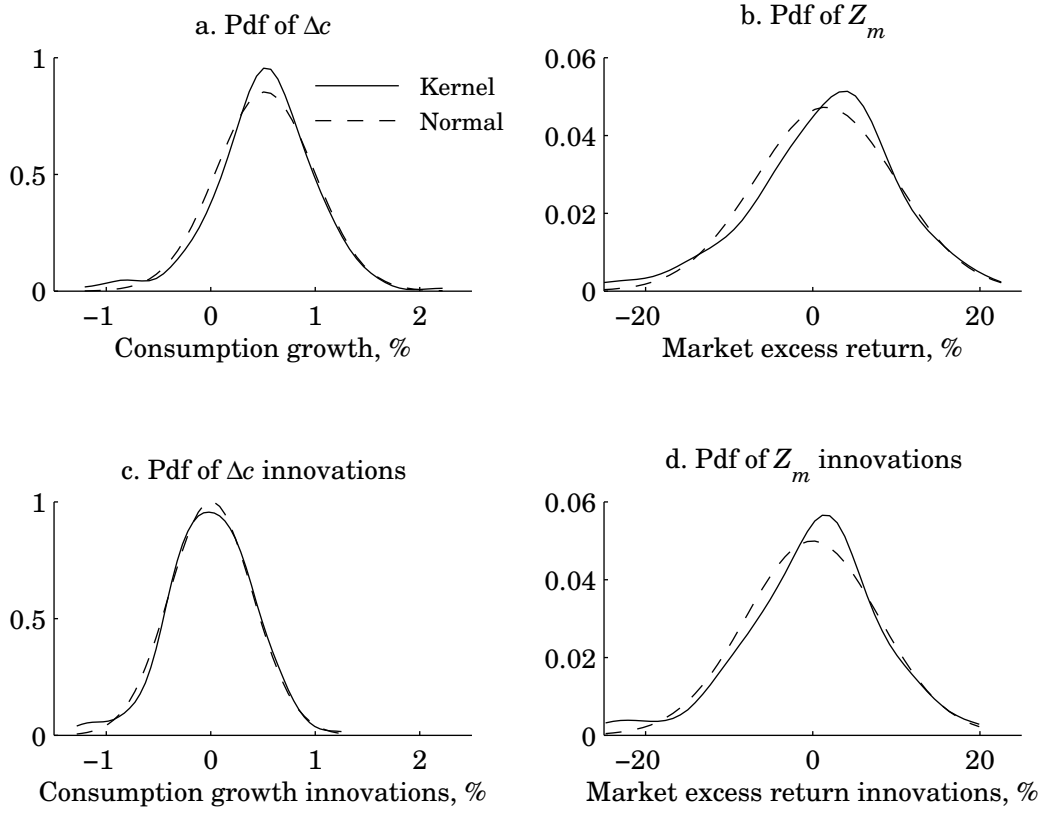
In any case, *Figure 1* suggests that these assumptions might be reasonable. The upper panel shows unconditional distributions of the growth of US real consumption per capita of nondurable goods and services and of the real excess return on a broad US equity index. (See Appendix A for details on the data). The non-parametric kernel density estimate of consumption growth is quite similar to a normal distribution, but this is not the case for

<sup>4</sup>For instance, (5) would be  $E(Z_{it}) = \beta_{ic} \text{Var}(\Delta c_t)\gamma$  where  $\beta_{ic} = \text{Cov}(Z_{it}, \Delta c_t) / \text{Var}(\Delta c_t)$ .

<sup>5</sup>Lettau and Uhlig (2002) also provide simple (but different) analytical expressions for these models—mostly with the aim to discuss the equity premium and riskfree rate puzzles.

<sup>6</sup>Otherwise, the relevant covariance would be between  $Z_{it}$  and  $(C_t/C_{t-1})^{-\gamma}$ .

the US market excess return which has a lot more skewness.



**Figure 1: Density functions of consumption growth and equity market excess returns, 1957–2004.** The kernel density function of a variable  $x$  is estimated by using a  $N(0, \sigma)$  kernel with  $\sigma = 1.06 \text{Std}(x)T^{-1/5}$ . The normal distribution is calculated from the estimated mean and variance of the same variable. See Appendix A for details on data sources and transformations.

While it is difficult to get information about the shape of the conditional distributions, the unconditional distribution of forecast errors may provide some clues, at least if the means of the conditional distributions change more than other moments. (It would be the correct measure if the other moments do not change at all.) The lower panel shows the distributions of forecast errors from a vector autoregression of consumption growth, equity market excess return, the “cay” variable of Lettau and Ludvigson (2001), and the 3-month T-bill rate. These distributions show the same general pattern as the unconditional distributions: consumption growth is almost normally distributed, but the market return is



not.<sup>7</sup>

## The Relation to the Equity Premium Puzzle

Most studies of the consumption-based asset pricing model have focused on the problem of reconciling the low volatility of consumption growth with the high excess return of an aggregate equity portfolio (this is the equity premium puzzle of Mehra and Prescott (1985)). The basic approach in those studies is to “fit” the unconditional pricing equation (5) to a broad equity market index only—and then assess whether the implied risk aversion coefficient is plausible.

To illustrate this approach, write (5) as

$$\underbrace{E(Z_{it})}_{6.5\%} = \underbrace{\text{Corr}(Z_{it}, \Delta c_t)}_{0.16} \underbrace{\text{Std}(Z_{it})}_{17\%} \underbrace{\text{Std}(\Delta c_t)}_{1\%} \gamma, \quad (7)$$

where the numbers are for the real excess return on US equity and consumption growth, 1957–2004. To fit this equation, the risk aversion,  $\gamma$ , needs to be around 250, which is an extremely high value. Even if we pretend that the correlation is as high as possible (1 instead of 0.16),  $\gamma$  needs to be around 40.

Recent asset pricing theory (discussed in detail below) has made substantial progress towards accounting for the equity premium puzzle by introducing recursive utility functions, habit persistence, or idiosyncratic risk. The way these new models achieve this is by either changing the relevant measure of consumption risk or the interpretation of the risk aversion parameter. The focus of this paper is to evaluate these models by studying their performance in a cross-section of returns.

### 2.3 C-CAPM Refinements

This section discusses several refinements of the C-CAPM: Epstein-Zin preferences, habit persistence, and idiosyncratic risk. It is shown that all these refinements imply (without adding too strong assumptions) that the conditional pricing model (6) should hold—although the risk aversion coefficient may need to be reinterpreted and might be time-

---

<sup>7</sup>In addition, implied (risk neutral) distributions derived from options on equity typically show distinct non-normal patterns, while the subjective distributions of output growth (as a proxy of consumption growth) reported by professional forecasters typically look similar to normal distributions—see Giordani and Söderlind (2005).

varying.

### The CRRA Model as a Special Case of Epstein-Zin Preferences

The recursive utility function in Epstein and Zin (1989) has different parameters for the risk aversion and the elasticity of intertemporal substitution. This generally gives a complicated pricing expression, but it coincides with the CRRA model in special cases. In particular, this happens if all wealth is marketable and there are no one-period innovations in the consumption-wealth ratio. To get this result we need either iid market returns or an elasticity of intertemporal substitution equal to one. (See Appendix B for details.) In these cases, testing the pricing performance of the CRRA model is the same as testing the Epstein-Zin model.

### Habit Persistence

The habit persistence model of Campbell and Cochrane (1999) has a CRRA utility function, but the argument is the difference between consumption and a habit level,  $C_t - X_t$ , instead of just consumption. The habit is parameterized in terms of the “surplus ratio”  $S_t = (C_t - X_t)/C_t$ , which measures how much aggregate consumption exceeds the habit. Since this ratio is external to the investor, marginal utility becomes  $(C_t - X_t)^{-\gamma} = (C_t S_t)^{-\gamma}$ . The log SDF is therefore

$$\ln M_t = -\gamma(\Delta s_t + \Delta c_t), \quad (8)$$

where  $s_t$  is the log surplus ratio. The process for  $s_t$  is assumed to be a non-linear AR(1) (constant suppressed)

$$s_t = \phi s_{t-1} + \lambda(s_{t-1})\Delta c_t, \quad (9)$$

where  $\lambda(s_{t-1}) \geq 0$  is a decreasing function of  $s_{t-1}$ .

It is straightforward to show (see Appendix B) that the conditional pricing expression is the same as (6), but with

$$\kappa_t = \gamma[1 + \lambda(s_{t-1})] \quad (10)$$

instead of just  $\gamma$ . The model therefore has the same cross-sectional implication as the conditional CRRA model. The only difference is that the slope coefficient will be time-varying since it involves  $1 + \lambda(s_{t-1})$ .

Since the log surplus ratio,  $s_t$ , is unobservable, the unconditional pricing expression

is not well suited for empirical testing—unless we make some further assumptions. In particular, it can be shown (see Appendix B) that if we assume that  $\lambda(s_{t-1})$  is a constant  $\lambda$  and that the excess return is unpredictable (at least by  $s_{t-1}$ ) then the unconditional pricing expression is the same as (5), but with

$$\kappa = \gamma(1 + \lambda) \quad (11)$$

instead of  $\gamma$ . The only difference to the standard CRRA model is the extra  $1 + \lambda$  term. This special case clearly ruins several of the interesting features of the model in Campbell and Cochrane (1999), so using (11) should probably be interpreted as focusing on a more standard habit persistence model. In contrast, the conditional expression (10) is a direct implication of the Campbell and Cochrane (1999) model.

### Idiosyncratic Risk

The volatility of aggregate consumption may underestimate the risk faced by investors if there are uninsurable individual shocks. Such shocks mean that the consumption growth of investor  $j$  is the aggregate consumption growth plus an idiosyncratic component. His log SDF is therefore  $\ln M_{jt} = -\gamma \Delta c_t - \gamma u_{jt+\delta}$ , where  $u_{jt+\delta}$  is the idiosyncratic shock. For analytical convenience, I let the shock be realized a split second ( $\delta$ ) after the asset return and aggregate consumption. Using the law of iterated expectations, the Euler equation of investor  $j$ ,  $E_{t-1}(Z_{it} M_{jt}) = 0$ , can now be written

$$E_{t-1}[Z_{it} \exp(-\gamma \Delta c_t) E_t \exp(-\gamma u_{jt+\delta})] = 0. \quad (12)$$

To simplify, assume that the distribution of the idiosyncratic shock, conditional on the information set in  $t$ , is normal. It is important that the mean of this distribution does not depend on the return (or else the shock is insurable) or consumption growth (or else the shock is non-idiosyncratic). I therefore assume that the mean is always zero. In contrast, the variance is assumed to be  $2\lambda(\varepsilon_t)$  where  $\lambda(\varepsilon_t)$  is some function of the aggregate shock,  $\varepsilon_t = \Delta c_t - E_{t-1} \Delta c_t$ . To simplify, I approximate  $\lambda(\varepsilon_t)$  by  $a + b_{t-1}\varepsilon_t$ , where  $b_{t-1}$  is known in  $t - 1$ .<sup>8</sup>

With these assumptions, it can be shown (see Appendix B) that the conditional pricing

---

<sup>8</sup>This is similar to Lettau (2002) who assumes that  $\varepsilon_t$  and  $\lambda(\varepsilon_t)$  have a bivariate normal distribution.

expression is the same as (6), but with

$$\kappa_t = \gamma(1 - \gamma b_{t-1}) \quad (13)$$

instead of  $\gamma$ . If  $b = 0$ , so idiosyncratic risk has a constant variance, then it drops out of the pricing expression and we are back in the standard CRRA model: idiosyncratic shocks have no effect on risk premia unless their variance depends on the aggregate shock (see Mankiw (1986) and Constantinides and Duffie (1996)).

Instead, if  $b_{t-1} < 0$ , so bad times (negative consumption surprise) are also risky times, then the idiosyncratic shocks add to the expected return. However, the implication for the cross-sectional pattern of expected returns is the same as in the conditional CRRA model without idiosyncratic shocks.

In contrast, the unconditional pricing expression is not that easily tested. However, if we assume that consumption growth is unpredictable and  $b_{t-1}$  is constant, then it can be shown (see Appendix B) that the unconditional pricing expression is the same as (5), but with  $\gamma(1 - \gamma b)$  instead of  $\gamma$ .

### Summary of the Testable Implications

All these models imply (without adding too strong assumptions) that the conditional pricing model (6) should hold—although the risk aversion coefficient may need to be reinterpreted and is likely to be time-varying (driven by habits and/or the movements in the variance of the idiosyncratic shocks).

With some additional strong assumptions (unpredictable consumption changes and excess returns, linear time series process for the log surplus ratio), also the unconditional pricing expression (5) should hold—provided we reinterpret the risk aversion coefficient.

## 3 Empirical Evidence

### 3.1 Unconditional Pricing

This section presents results for unconditional average excess returns and covariances with consumption growth.

## The Estimation Framework

To estimate and test the model, a GMM framework is used. The moment conditions give the traditional estimators of the means and covariances, and the slope coefficient in (5) is estimated as a LS regression of the mean returns on the covariances (and a constant).

Allowing for a constant implies that we focus on how differences in covariances are related to differences in risk premia. That is, the focus is on whether C-CAPM can explain the variation in risk premia among the 25 portfolios, not the general level of them.

By estimating this system jointly, the uncertainty about the means and covariances are incorporated into the test of the slope coefficient and errors in (5). See Appendix C for details.

The asset data are real quarterly returns on the 25 Fama and French (1993) equally weighted portfolios in excess of a real T-bill rate. The consumer price index is used as the deflator for the returns. The real consumption per capita is for nondurable goods and services, and is measured leading one quarter.<sup>9</sup> More details on data are given in Appendix A.

The historical (1957–2004) average excess returns (“risk premia”) are shown in *Table 1*. The table forms the portfolios into a  $5 \times 5$  matrix where the firms in cell  $ij$  belong to the  $i$ th quintile of firm size and  $j$ th quintile of book-to-market value ratio (B/M). Smaller firms have had higher risk premia than larger firms, and value firms (high book-to-market ratios) have had higher risk premia than growth firms (low book-to-market ratios). There is therefore a non-trivial cross-section of risk premia for the model to explain.

## The Estimation Results

The main result is that the unconditional model is soundly rejected. The hypothesis that all the pricing errors (deviations from (5)) are zero, is rejected at a significance level below 1%. For comparison, it can be noticed that the rejection of the traditional CAPM is as strong.<sup>10</sup> Although this finding is based on a traditional (asymptotic) GMM approach, it is also verified by a Monte Carlo study: the simulations suggest that the model should

---

<sup>9</sup>It is unclear if returns should be related to what is recorded as consumption this quarter or the next. The reason is that consumption is measured as a flow during the quarter, while returns are measured at the end of the quarter. I choose to show the results for  $\Delta c_{t+1}$  since they are somewhat more supportive of the consumption-based model—but the difference is not large.

<sup>10</sup>To derive CAPM from (3), assume that the log SDF is an affine function of the market excess return.

	B/M				
	1	2	3	4	5
Size 1	7.6	12.4	13.5	15.7	18.6
2	5.4	9.2	11.7	12.1	13.5
3	5.9	9.4	9.2	11.5	12.8
4	6.9	7.2	9.8	10.7	11.2
5	5.8	7.0	7.6	7.4	8.4

**Table 1: Historical risk premia (annualised %), 1957–2004.** Results are shown for the 25 equally-weighted Fama-French portfolios, formed according to size and book-to-market ratios (B/M).

be rejected at a somewhat higher significance level—but still below 1%. (See Section 3.2 for details on the simulation experiment.)


*Figure 2* illustrates the findings. The vertical axis shows (annualised) mean excess returns and the horizontal axis shows unconditional covariances with consumption growth. The data points and the fitted line from the estimation are marked, although the scales of the axes are different than used in the estimation (the figure has means that are scaled by 100 and covariances that are scaled by 10,000). The overall pattern is that a higher covariance is associated with a higher risk premium (the  $R^2$  is 0.4), but there are large deviations.

The “pricing errors” (difference between historical and fitted risk premia) are shown in *Table 2*. There is a great dispersion (around zero), which is the reason for the rejection of the model. The results are dominated by the tendency of “mean reversion” of the model: portfolios with high historical risk premia (see *Table 1*) are underestimated (have positive pricing errors) and vice versa. This tendency is natural (and also obvious from *Figure 2*), but the main concern is the magnitude of the errors: the pricing errors appear to be very large. This impression is emphasised by the relative pricing errors (errors divided by the historical risk premia) shown in *Table 3*. The model completely fails in pricing the growth firms (low B/M), whereas the other portfolios seem to be more in line with the model predictions.

In general, these results are consistent with earlier findings in, for instance, Breeden, Gibbons, and Litzenberger (1989) and Lettau and Ludvigson (2001).<sup>11</sup>

The point estimate of the slope coefficient is above 200 (see *Figure 2*), which is indeed

<sup>11</sup> Bansal, Dittmar, and Lundblad (2005) get more positive results by focusing on the covariance between consumption and cash flows.



**Figure 2: Expected excess returns and the covariance with consumption growth (unconditional), 1957–2004.** Excess returns are real returns in excess of the real return on a one-month Treasury Bill. To annualise the expected returns and covariances, quarterly figures are multiplied by 4. See Appendix A for details on data sources and transformations.

	B/M				
	1	2	3	4	5
Size 1	−6.7	−1.6	1.0	3.2	4.7
2	−4.0	−0.6	1.8	2.1	1.7
3	−3.7	0.8	0.1	2.0	2.7
4	−1.6	−1.3	1.7	2.1	−0.2
5	−3.0	−0.2	0.2	−0.3	−0.9

**Table 2: Historical minus fitted risk premia (annualised %) from the unconditional model, 1957–2004.** Results are shown for the 25 equally-weighted Fama-French portfolios, formed according to size and book-to-market ratios (B/M).

very high. In any case, we should not interpret this as a direct measure of the traditional risk aversion coefficient—since it may (under the strong assumptions discussed above) reflect risk aversion as well as habit persistence and idiosyncratic risk. The t-statistic of the slope coefficient is 1.7, so the hypothesis of a zero slope can be rejected at the 10% level, and the hypothesis of a negative slope can be rejected at the 5% level.

To sum up, the bulk of the evidence on the unconditional model is against the basic CRRA model—mostly because it does not price the 25 assets with sufficient precision.

	B/M				
	1	2	3	4	5
Size 1	−88.7	−13.2	7.1	20.5	25.2
2	−74.4	−6.1	15.4	17.0	12.5
3	−63.1	8.3	1.6	17.5	21.4
4	−22.9	−18.1	17.8	20.1	−1.8
5	−52.7	−3.2	3.1	−3.5	−10.9

**Table 3: Relative errors of risk premia (in %) of the unconditional model, 1957–2004.** The relative errors are defined as historical minus fitted risk premia, divided by historical risk premia. Results are shown for the 25 equally-weighted Fama-French portfolios, formed according to size and book-to-market ratios (B/M).

However, this can only be interpreted as evidence against the refined models (habit persistence and idiosyncratic risk) under very strong assumptions. To say more about the refinements, we need to study the conditional pricing equation.

### 3.2 Conditional Pricing

This section presents results for conditional (time-varying) risk premia and covariances.

#### The Estimation Framework

To estimate the conditional moments, I use a combination of a VAR model and a multivariate volatility model. The models are estimated in a recursive way (longer and longer sample), and all moments (expected values and covariances) are based on “out-of-sample” forecasts. This approach retains the well-known efficiency of VAR models, while avoiding the risk of “in-sample” overfitting.<sup>12</sup> To be precise, I estimate the following model.

*First*, I specify a simple VAR(2) model with four variables: consumption growth, the excess return on the US equity market ( $R_{mt}^e$ ), the “cay” variable of Lettau and Ludvigson (2001) ( $cay_t$ ) and a 3-month T-bill rate ( $r_t$ ). Let  $x_t$  be the vector  $[\Delta c_t, R_{mt}^e, cay_t, r_t]$ . The VAR is estimated recursively (longer and longer sample). The first estimation is done with data for 1947Q1–1956Q4 and a prediction is then made for 1957Q1—the difference to the actual outcome is then the out-of-sample forecast error for 1957Q1. The sample is

<sup>12</sup>See Goyal and Welch (2004) and Campbell and Thompson (2005) for a recent discussion of the in-sample versus out-of-sample predictability of equity returns.



then extended one quarter, and the forecast error for 1957Q2 is generated, and so forth. The forecasting equation for consumption growth in  $t$  is then an affine function of the vectors  $x_{t-1}$  and  $x_{t-2}$ . If we let  $u_t$  be the forecast error, we can write actual consumption growth as

$$\Delta c_t = a_{0t-1} + a_{1t-1}x_{t-1} + a_{2t-1}x_{t-2} + u_t, \quad (14)$$

where the coefficients ( $a_{it-1}$ ) carry time subscripts to remind us that they are estimated on data up to and including  $t - 1$ .

*Second*, the conditional asset pricing equation

$$Z_{it} = \alpha_t + \kappa_t \omega_{it-1} + \varepsilon_{it} \quad (15)$$

is estimated period-by-period—by using the cross-section of returns and conditional covariances (denoted  $\omega_{it-1}$ , see below). This allows both the intercept and the slope coefficient to vary freely between periods.

*Third* (and finally), the conditional covariance (of asset  $i$  and consumption growth) is estimated by an exponentially moving average estimator similar to the RiskMetrics approach

$$\omega_{it-1} = 0.25\varepsilon_{it-1}u_{t-1} + 0.75\omega_{it-2}. \quad (16)$$

The last two equations form a kind of “GARCH-in-mean” system, where the conditional covariance (known in  $t - 1$ ) is allowed to predict the return in period  $t$ —just as suggested by the conditional pricing model. Replacing the RiskMetrics approach with a “dynamic conditional correlation multivariate GARCH” (see Engle (2002)) gives similar results, but with occasional numerical problems (especially in the Monte Carlo experiment discussed below), which speaks in favour of the simpler approach. The weights of 0.25 and 0.75 are approximately the same as the corresponding coefficients in the GARCH model.

This approach is designed to focus on whether the cross-sectional variation in covariances can explain the cross-sectional variation in returns (just like the unconditional test). This means that there are some special features that need explaining. Consider the cross-sectional regression (15) for period  $t - 1$ . It has actual returns (not expected returns) as the dependent variable. Therefore, the time-varying intercept ( $\alpha_{t-1}$ ) picks up common movements (predictable or not) in returns, and the residual ( $\varepsilon_{it-1}$ ) will pick up asset specific (idiosyncratic) news about returns, but no common innovations. As a consequence, the “conditional covariance”  $\omega_{it-1}$  in (16)—which is based on the residuals ( $\varepsilon_{it-1}$ ) from

the cross-sectional regression—leaves out one component of the true covariance: the innovations in returns that are common for all assets. Of course, that does not matter for the results of the cross-sectional regression in  $t$  (when  $\omega_{it-1}$  is a regressor), since common movements in covariances are anyway loaded into the intercept ( $\alpha_t$ ).

The estimation framework used here differs considerably from the scaled factor model in Lettau and Ludvigson (2001). They impose the restriction that the time variation (using a beta representation) is a linear function of some conditioning variables (specifically, the cay variable) only—and they cannot distinguish between the time variation in the effective risk aversion and the time variation in the covariance. In contrast, (15) and (16) allow the effective (or local) risk aversion to vary freely between periods at the same time as an explicit measure of the conditional covariance is used: this allows a clean interpretation of the estimated slope coefficient.

Rather, the estimation framework is more similar to that in Duffee (2005), who also estimates a conditional covariance—by projecting the cross-product of the innovations ( $\varepsilon_{it-1}u_{t-1}$  in (16)) on to a set of instruments. Overall, the main difference is that the current paper focuses on the cross-section of returns, whereas Duffee (2005) focused on the market returns. (A somewhat smaller difference is that Duffee (2005) restricts the effective risk aversion to be a time-invariant function of some prespecified macro-variables.)


The potential drawbacks of the formulation in (15) are that the cross-section (here 25 portfolios) may be too small, so the estimates of the slope coefficient become noisy—and that the model for the conditional covariance may be misspecified.

## The Estimation Results

The main result is that the conditional model is easily rejected. The hypothesis that all assets are (on average, over time) correctly priced is rejected at a significance level below 1%. That is, the time-averages of the pricing error of the model (15),  $\sum_{t=1}^T \varepsilon_{it}/T$ , are significantly different from zero for at least some assets. Since the model cannot get the time-average of the excess returns right, there is little point in looking at more stringent tests (for instance, whether the pricing errors are forecastable).

The test is done by forming a Wald statistic, where the covariance matrix of the average pricing errors is estimated by a Newey-West method (using one lag). This approach is complicated by the fact that the cross-sectional regression contains “generated regres-

sors.” It would perhaps be possible to set up a GMM framework to handle this (as I did for the unconditional test), but the time-varying parameters would require a very large number of moment conditions—and the properties of this system is virtually unknown. Instead, I perform a Monte Carlo simulation based on (14)–(16) with 25,000 replications. See Appendix C for details. The simulated distribution looks like a  $\chi^2$  distribution (as we may have hoped) but with more degrees of freedom than expected—the distribution is moved to the right. The same holds for the unconditional test, but to a smaller extent. In any case, the hypothesis that all assets have zero time-averages of the pricing errors is still easily rejected at the 1% significance level.



**Figure 3: Average excess returns (over and above the market wide return) and fitted values from the conditional model, 1957–2004.** The figure shows  $\Sigma_{t=1}^T (Z_{it} - \alpha_t) / T = \Sigma_{t=1}^T (\kappa_t \omega_{it-1} + \varepsilon_{it}) / T$  from (15) against  $\Sigma_{t=1}^T \kappa_t \omega_{it-1} / T$ .

Figure 3 illustrates the unconditional performance of the conditional model. In terms of (15), the figure shows the averages of  $\kappa_t \omega_{it-1} + \varepsilon_{it}$  against the averages of their fitted values ( $\kappa_t \omega_{it-1}$ ). This is a much more revealing (and fairer) evaluation than just comparing the average returns—since the model in (15) has a free parameter ( $\alpha_t$ ) for the market-wide return in every period: the figure highlights the contribution of the conditional covariance. Overall, the conditional model appears to do approximately as well as the unconditional model, or even slightly worse (the uncentred  $R^2$  is only 0.33, compared to 0.40 for the unconditional model).

Table 4 shows the (time-averages of the) pricing errors. The pattern is fairly similar to the results from the unconditional model: the pricing errors are of similar magnitudes, and there is a general tendency of mispricing all growth firms and also the small value firms.

	B/M				
	1	2	3	4	5
Size 1	−2.0	1.2	2.9	4.4	7.0
2	−3.8	−1.5	1.3	0.9	1.7
3	−3.2	−0.4	−1.6	1.5	2.2
4	−1.6	−2.1	0.4	1.2	−0.3
5	−2.5	−1.4	−0.9	−1.8	−1.5

**Table 4: Historical minus fitted average risk premia (annualised %) from the conditional model, 1957–2004.** Results are shown for the 25 equally-weighted Fama-French portfolios, formed according to size and book-to-market ratios (B/M).

It is thus clear that the model fails the econometric tests. But, could it still make some sense—from an economic point of view? That is, do the estimated slope coefficient (the effective or local risk aversion) behave in a way that is consistent with the refinements of C-CAPM? Several pieces of evidence suggest not.


*First*, the slope coefficients in (15) are often estimated to be negative. A Bonferroni joint test of the hypothesis that all the slope coefficients are positive<sup>13</sup> is rejected using both traditional and simulated critical values at the 10% level (but not at the 5% level). To reduce sampling uncertainty (noise) Figure 4 shows (centred) moving averages of  $\pm 2$  quarters of the slope coefficients. There are clearly many periods of negative values—which is at odds with the theoretical models surveyed above.

*Second*, the estimated coefficients fluctuate a lot—and in a way that appears hard to reconcile with the ideas behind the C-CAPM refinements. There is some relation to the NBER recessions (shaded areas in Figure 4)—the slope coefficient (effective or local risk aversion) increases rapidly during several recessions. On the other hand, it fails to increase during several other recessions, and there are many dramatic movements that appear to be unrelated to the business cycle. Correlations with the macro variables in the

<sup>13</sup>The test at the 5% significance level compares the smallest t-statistic with the 0.05/(number of slope coefficients) percentile of a standard normal distribution. See, for instance, Mittelhammer, Judge, and Miller (2000).

VAR discussed before are very low (at most 0.06), and the same is true for the correlation with the equity market volatility. The only possible exception is that there is a positive correlation (around 0.15) between the slope coefficient and a moving average of the equity market return.

Since asset pricing is all about beliefs about the future, it is also interesting to relate the results to survey data. The Survey of Professional Forecasters is a quarterly survey of some 30 forecasters' views on key economic variables, conducted by the Federal Reserve Bank of Philadelphia. The respondents, who supply anonymous answers, are professional forecasters from the business and financial community. (See Appendix A for details on the data.) It turns out that there is little correlation between the slope coefficient and the (median) point forecast of future inflation, GDP growth, or business profitability. The best support for the model is that there is a reasonably strong correlation with uncertainty about future inflation (1969–2004) and also uncertainty about future GDP growth (1982–2004), but this evidence is only from relatively short subsamples.



**Figure 4: Slope coefficient in conditional pricing equation, 1957–2004.** The figure shows a centred moving average of  $\pm 2$  quarters of the slope coefficient in (15), estimated on the 25 equally-weighted Fama-French portfolios, formed according to size and book-to-market ratios (B/M).

To sum up, the conditional model gives little support to the C-CAPM refinements: the pricing errors are large and the estimated effective risk aversions do not appear to be correlated with the sort of macro variables that the C-CAPM refinements suggest.

These findings are quite different from those in Lettau and Ludvigson (2001) who

find that a scaled factor model works reasonably well. This suggests that it might be important to have an explicit measure of the conditional covariance. The findings are somewhat more in line with those of Duffee (2005), who showed that there is a negative relation between the conditional covariance and the conditional expected market return. Although the current paper focuses on the cross-sectional aspects, it also finds a negative slope coefficient for many periods (and several other problems with the model).

## 4 Summary

This paper studies if the consumption-based asset pricing model can explain the cross-sectional dispersion of conditional expected returns.

It is shown that the CRRA model and friends (Epstein-Zin utility, habit persistence, and idiosyncratic shocks) share the same implications for the cross-sectional dispersion of risk premia: the conditional expected return should be linearly increasing in the conditional covariance of the asset return with aggregate consumption growth—but with a time-varying slope coefficient.

This is studied on quarterly data for the 25 Fama-French portfolios (1957–2004). There is little support for the consumption-based model: the pricing errors are large and the slope coefficient (the effective risk aversion) seems to be unrelated to most variables that could proxy for idiosyncratic risks or habits.

## A Data Appendix

The nominal *stock returns* and the *nominal interest rate* are from the web site of French (2001). These monthly returns are converted to quarterly returns by multiplying the monthly gross returns, for instance, the gross returns for January, February, and March are multiplied to generate a quarterly gross return. The 25 (equally weighted) portfolios are formed along the quintiles of size (firm market values) and book value/market values.

*Real returns* are calculated by dividing the nominal gross return by the gross inflation rate over the same period. Inflation is calculated from the seasonally adjusted CPI for all urban consumers (available at <http://research.stlouisfed.org/fred2/>).

Quarterly growth of *real consumption per capita* of nondurable goods and services is calculated from the seasonally adjusted number in NIPA tables (available at <http://www.bea.doc.gov/>). The growth rate is calculated as a weighted average of the growth rate of nondurable goods and the growth rate of services (chained 2000 dollars), where the (time-varying) weight is the relative (current dollar) size of nondurable goods in relation to services.

The data of the Survey of Professional Forecasters is from the Federal Reserve Bank of Philadelphia (<http://www.phil.frb.org/>). For details, see Croushore (1993). Also, see Giordani and Söderlind (2005) for a description of how to estimate the uncertainty from the probability distributions in the survey.

## B Asset Pricing Appendix

### B.1 Derivation of Stein's Lemma for a Special Case of Mixture Normals

This section proves that Stein's lemma continues to hold if  $x$  and  $y$  have a bivariate mixture normal distribution, provided the marginal distribution of  $y$  is normal.

Let the pdf of  $(x, y)$  be a mixture of  $n$  bivariate normal distributions

$$pdf \left( \begin{bmatrix} x \\ y \end{bmatrix} \right) = \sum_{i=1}^n \alpha_i \phi \left( \begin{bmatrix} x \\ y \end{bmatrix}; \begin{bmatrix} E_i(x) \\ E_i(y) \end{bmatrix}, \begin{bmatrix} \text{Var}_i(x) & \text{Cov}_i(x, y) \\ \text{Cov}_i(x, y) & \text{Var}_i(y) \end{bmatrix} \right), \text{ with } \sum_{i=1}^n \alpha_i = 1,$$

and where  $\phi(z; \mu, \Sigma)$  is the normal pdf with mean vector  $\mu$  and covariance matrix  $\Sigma$ .

Direct calculations give

$$\text{Cov}[x, h(y)] = \sum_{i=1}^n \alpha_i \{ \text{Cov}_i[x, h(y)] + E_i(x) E_i[h(y)] \} - E(x) E[h(y)].$$

If  $E_i[h(y)]$  is a constant  $E[h(y)]$ , then this simplifies to

$$\text{Cov}[x, h(y)] = \sum_{i=1}^n \alpha_i \text{Cov}_i[x, h(y)].$$

Since  $\text{Cov}_i[x, h(y)] = \text{Cov}_i(x, y) E_i[h'(y)]$  (Stein's lemma) for all states, we get (since  $E_i[h'(y)] = E[h'(y)]$ )

$$\text{Cov}[x, h(y)] = E[h'(y)] \sum_{i=1}^n \alpha_i \text{Cov}_i(x, y).$$

If  $E_i(y) = E(y)$ , then the sum equals  $\text{Cov}(x, y)$ , following the same argument as above. Note, however, that for both  $E_i(y) = E(y)$  and  $E_i[h(y)] = E[h(y)]$  to be true requires that  $\text{Var}_i(y) = \text{Var}(y)$ . This means that the marginal distribution of  $y$  must be a normal distribution.

## B.2 Derivation of (2) and (3) (Basic Model)

If  $x$  and  $y$  have a bivariate normal distribution and  $h(y)$  is a differentiable function such that  $E[|h'(y)|] < \infty$ , then  $\text{Cov}[x, h(y)] = \text{Cov}(x, y) E[h'(y)]$ . The moments can be interpreted as either unconditional or conditional moments. See, for instance, Cochrane (2001). Here, this gives  $\text{Cov}(Z_{it}, M_t) = \text{Cov}(Z_{it}, \ln M_t) E(M_t)$ . Since (1) can be written  $E(Z_{it}) E(M_t) = -\text{Cov}(Z_{it}, M_t)$ , (2) and (3) follow.

## B.3 Derivation of the Results on Epstein-Zin Utility

Epstein and Zin (1989) show that if all wealth is marketable, then the Euler equation for the excess return of asset  $i$  is

$$E_{t-1}[(C_t/C_{t-1})^{-\theta/\psi} R_{mt}^{\theta-1} Z_{it}] = 0, \text{ where } \theta = (1 - \gamma)/(1 - 1/\psi),$$

where  $R_{mt}$  is the market gross return,  $\gamma$  the risk aversion, and  $\psi$  the elasticity of intertemporal substitution.

If  $C_t/W_t = 1/\alpha_t$ , then we can substitute for wealth in the budget restriction to get  $C_t \alpha_t = R_{mt}(\alpha_{t-1} C_{t-1} - C_{t-1})$ , or  $C_t/C_{t-1} = R_{mt}(\alpha_{t-1} - 1)/\alpha_t$  which in turn gives  $(C_t/C_{t-1})\alpha_t/(\alpha_{t-1} - 1) = R_{mt}$ . Using in the Euler equation gives  $E_{t-1}[(C_t/C_{t-1})^{-\gamma} \alpha_t^{\theta-1} Z_{it}] = 0$ . Campbell (1993) shows that there are no innovations in  $\alpha_t$  if  $\psi = 1$  and that  $\alpha$  is a constant if the market returns are iid or if  $\psi = 1$  and all innovations are homoskedastic.

## B.4 Derivation of (10) and (11) (Habit Persistence)

Use (9) in (8) to get  $\ln M_t = -\gamma(\phi - 1)s_{t-1} - [1 + \lambda(s_{t-1})]\gamma \Delta c_t$ . Since  $s_{t-1}$  is known in  $t-1$ , the conditional covariance is  $\text{Cov}_{t-1}(Z_{it}, \ln M_t) = \text{Cov}_{t-1}(Z_{it}, \Delta c_t)[1 + \lambda(s_{t-1})]\gamma$ . Use in (2) to get (10).

If  $\lambda(s_{t-1})$  is a constant  $\lambda$ , then unconditional equation (3) becomes  $E(Z_{it}) = \text{Cov}[Z_{it}, \gamma(\phi - 1)s_{t-1} + (1 + \lambda)\gamma \Delta c_t]$ . The  $s_{t-1}$  term cancels if it cannot predict  $Z_{it}$  which gives  $E(Z_{it}) = \text{Cov}(Z_{it}, \Delta c_t) \gamma(1 + \lambda)$ .



## B.5 Derivation of (13) (Idiosyncratic Risk)

We have that  $E_t \exp(-\gamma u_{jt+\delta}) = \exp[\gamma^2 \lambda(\varepsilon_t)]$ , and  $\Delta c_t = E_{t-1} \Delta c_t + \varepsilon_t$ . Equation (12) can therefore be written  $E_{t-1}\{Z_{it} \exp[-\gamma E_{t-1} \Delta c_t - \gamma \varepsilon_t + \gamma^2 \lambda(\varepsilon_t)]\} = 0$ . Letting  $\lambda(\varepsilon_t) = a + b_{t-1} \varepsilon_t$  and cancelling the non-random term ( $\gamma^2 a$ ) gives  $E_{t-1}\{Z_{it} \exp[-\gamma E_{t-1} \Delta c_t - \varepsilon_t \gamma (1 - b_{t-1} \gamma)]\} = 0$ . Clearly, the  $\exp[\cdot]$  term then corresponds to  $M_t$  in (1).

The conditional expression (2) becomes

$$\begin{aligned} E_{t-1}(Z_{it}) &= -\text{Cov}_{t-1}[Z_{it}, -\gamma E_{t-1} \Delta c_t - \gamma \varepsilon_t (1 - b_{t-1} \gamma)] \\ &= \text{Cov}_{t-1}(Z_{it}, \varepsilon_t) \gamma (1 - b_{t-1} \gamma), \end{aligned}$$

since  $E_{t-1} \Delta c_t$  is known already in  $t - 1$ . We can replace  $\varepsilon_t$  in this expression by  $\Delta c_t$  if we want to.

The unconditional expression (3) becomes

$$E(Z_{it}) = \text{Cov}[Z_{it}, \gamma E_{t-1} \Delta c_t + \varepsilon_t \gamma (1 - b_{t-1} \gamma)].$$

If  $E_{t-1} \Delta c_t$  cannot predict  $Z_{it}$ , then  $\gamma E_{t-1} \Delta c_t$  cancels. If, in addition,  $E_{t-1} \Delta c_t$  and  $b_{t-1}$  are constants (which certainly implies that they cannot forecast  $Z_{it}$ ), then we get

$$E(Z_{it}) = \text{Cov}(Z_{it}, \Delta c_t) \gamma (1 - b \gamma).$$

## C Econometric Appendix

### C.1 Unconditional Test

Let there be  $N$  assets. The original moment conditions are

$$g_T(\beta) = \frac{1}{T} \sum_{t=1}^T \begin{bmatrix} (\Delta c_t - \mu_{\Delta c}) = 0 \\ (Z_{it} - \mu_i) = 0 \text{ for } i = 1, 2, \dots, N \\ [(\Delta c_t - \mu_c)(Z_{it} - \mu_i) - \sigma_{ci}] = 0 \text{ for } i = 1, 2, \dots, N \\ (Z_{it} - \alpha - \sigma_{ci} \kappa) = 0 \text{ for } i = 1, 2, \dots, N, \end{bmatrix}$$

where  $\mu_{\Delta c}$  is the mean of  $\Delta c_t$ ,  $\mu_i$  the mean of  $Z_{it}$ ,  $\sigma_{ci}$  the covariance of  $\Delta c_t$  and  $Z_{it}$ . This gives  $1 + 3N$  moment conditions and  $2N + 3$  parameters, so there are  $N - 2$  overidentifying restrictions.

To estimate, we define the combined moment conditions as

$$\begin{aligned} Ag_T(\beta) &= \mathbf{0}_{(2N+3) \times 1}, \text{ where} \\ A_{(2N+3) \times (1+3N)} &= \begin{bmatrix} 1 & \mathbf{0}_{1 \times N} & \mathbf{0}_{1 \times N} & \mathbf{0}_{1 \times N} \\ \mathbf{0}_{N \times 1} & I_N & \mathbf{0}_{N \times N} & \mathbf{0}_{N \times N} \\ \mathbf{0}_{N \times 1} & \mathbf{0}_{N \times N} & I_N & \mathbf{0}_{N \times N} \\ 0 & \mathbf{0}_{1 \times N} & \mathbf{0}_{1 \times N} & \sigma'_{ic} \\ 0 & \mathbf{0}_{1 \times N} & \mathbf{0}_{1 \times N} & \mathbf{1}_{1 \times N} \end{bmatrix}. \end{aligned}$$

These moment conditions mean that means and covariances are estimated in the tradi-

tional way, and that  $\kappa$  is estimated by a LS regression of  $E Z_{it}$  on a constant and  $\sigma_{ci}$ . The test that the pricing errors are all zero is a Wald test that  $g_T(\beta)$  are all zero, where the covariance matrix of the moments are estimated by a Newey-West method (using one lag).

See Cochrane (2001) 11 for the formulas for the covariance matrix of the parameters and the test of the sample moment conditions.

## C.2 Monte Carlo Simulations

The data generating process in the simulations of (14)–(16) is as follows.

Consumption growth is generated by simulating the VAR discussed in connection with (14): the VAR parameters are those estimated from the full historical sample (not time-varying), and the covariance matrices of the shocks (time-varying) are from an adjusted RiskMetrics estimation based on the historical residuals.

The adjustment amounts to making sure that the covariance matrix is guaranteed to be positive definite. This is achieved by an approach similar to that in models of “dynamic conditional correlation multivariate GARCH” (see Engle (2002)). In short, the procedure is this: (i) first the (time-varying) variances are estimated; (ii) (time-varying) covariance matrices of standardized residuals are estimated; (iii) the (time-varying) correlation matrices of the standardized residuals from step (ii) are recombined with the (time-varying) variances from step (i) to create time-varying covariances matrices.

The returns are generated from (15)–(16) where all the parameters are time-varying and taken from historical estimations, including the (time-varying) covariance matrices of the residuals in (15)—which are estimated by the adjusted RiskMetrics approach discussed above.

To do the simulations of the unconditional test, I scramble the historical  $\kappa_t$  and  $\omega_{it-1}$  so they are uncorrelated. The simulated data then satisfy an unconditional model, since by taking unconditional expectations of (15) we have  $E Z_{it} = E \alpha_t + E \kappa_t E \omega_{it-1}$ .

To simulate the properties of the Bonferroni test, all  $\kappa_t$  are set to zero.

## References

- Bansal, R., R. F. Dittmar, and C. T. Lundblad, 2005, "Consumption, Dividend, and the Cross Section of Equity Returns," *Journal of Finance*, 60, 1639–1672.
- Breeden, D. T., M. R. Gibbons, and R. H. Litzenberger, 1989, "Empirical Tests of the Consumption-Oriented CAPM," *Journal of Finance*, 44, 231–262.
- Campbell, J. Y., 1993, "Intertemporal Asset Pricing without Consumption Data," *American Economic Review*, 83, 487–512.
- Campbell, J. Y., 2001, "Consumption-Based Asset Pricing," in George Constantinides, Milton Harris, and Rene Stultz (ed.), *Handbook of the Economics of Finance*, North-Holland, Amsterdam.
- 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 S. B. Thompson, 2005, "Predicting the Equity Premium Out of Sample: Can Anything Beat the Historical Average," NBER WP 11468.
- Cochrane, J. H., 2001, *Asset Pricing*, Princeton University Press, Princeton, New Jersey.
- Constantinides, G. M., and D. Duffie, 1996, "Asset Pricing with Heterogeneous Consumers," *The Journal of Political Economy*, 104, 219–240.
- Croushore, D., 1993, "Introducing: The Survey of Professional Forecasters," *Federal Reserve Bank of Philadelphia Business Review*, pp. 3–13.
- Duffee, G. R., 2005, "Time Variation in the Covariance between Stock Returns and Consumption Growth," *Journal of Finance*, 60, 1673–1712.
- Engle, R. F., 2002, "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models," *Journal of Business and Economic Statistics*, 20, 339–351.

- Epstein, L. G., and S. E. Zin, 1989, "Substitution, Risk Aversion, and the Temporal Behavior of Asset Returns: An Empirical Analysis," *Journal of Political Economy*, 99, 263–286.
- Fama, E. F., and K. R. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.
- French, K. R., 2001, "U.S. Research Returns Data," <http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/>.
- Giordani, P., and P. Söderlind, 2005, "Is there Evidence of Pessimism and Doubt in Subjective Distributions? Implications for the Equity Premium Puzzle," *Journal of Economic Dynamics and Control*, forthcoming.
- Goyal, A., and I. Welch, 2004, "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction," NBER WP 10483.
- Lettau, M., 2002, "Idiosyncratic Risk and Volatility Bounds, or, Can Models with Idiosyncratic Risk Solve the Equity Premium Puzzle?," *Review of Economics and Statistics*, 84, 376–380.
- Lettau, M., and S. Ludvigson, 2001, "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying," *Journal of Political Economy*, 109, 1238–1287.
- Lettau, M., and H. Uhlig, 2002, "Sharpe Ratios and Preferences: An Analytical Approach," *Macroeconomic Dynamics*, 6, 242–265.
- Mankiw, G. N., 1986, "The Equity Premium and the Concentration of Aggregate Shocks," *Journal of Financial Economics*, 17, 211–219.
- Mehra, R., and E. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161.
- Mittelhammer, R. C., G. J. Judge, and D. J. Miller, 2000, *Econometric Foundations*, Cambridge University Press, Cambridge.