CAY Revisited:
Can Optimal Scaling Resurrect the (C)CAPM?[1]

Devraj Basu
Alexander Stremme

Warwick Business School, University of Warwick

November 2005

Currently under Review; Please do not Quote or Distribute

address for correspondence:

Alexander Stremme
Wariwck Business School
University of Warwick
Coventry, CV4 7AL
United Kingdom

e-mail: `alex.stremme@wbs.ac.uk`

[1] We thank Wayne Ferson for helpful discussions. We also thank Kenneth French and Sydney Ludvigson for making available the data used in this study. All remaining errors are ours.

CAY Revisited:
Can Optimal Scaling Resurrect the (C)CAPM?

May 2005

Abstract

In this paper, we evaluate specification and pricing error for the Consumption (C-) CAPM in the case where the model is optimally scaled by consumption-wealth ratio (CAY). Lettau and Ludvigson (2001b) show that the C-CAPM successfully explains a large portion (about 70%) of the cross-section of expected returns on Fama and French's size and book-to-market portfolios, when the model is scaled *linearly* by CAY. In contrast, we use the methodology developed in Basu and Stremme (2005) to construct the *optimal* factor scaling as a (possibly non-linear) function of the conditioning variable (CAY), designed to minimize the model's pricing error. We use a new measure of specification error, also developed in Basu and Stremme (2005), which allows us to analyze the performance of the model both in and out-of-sample.

We find that the optimal factor loadings are indeed non-linear in the instrument, in contrast to the linear specification prevalent in the literature. While our optimally scaled C-CAPM explains about 80% of the cross-section of expected returns on the size and book-to-market portfolios (thus in fact out-performing the linearly scaled model of Lettau and Ludvigson (2001b)), it fails to explain the returns on portfolios sorted by industry. Moreover, although the optimal use of CAY does dramatically improve the performance of the model, even the scaled model fails our specification test (for either set of base assets), implying that the model still has large pricing errors. Out-of-sample, the performance of the model deteriorates further, failing even to explain any significant portion of the cross-section of expected returns. For comparison, we also test a scaled version of the classic CAPM and find that it has in fact smaller pricing errors than the scaled C-CAPM.

JEL Classification: C31, C32, G11, G12

Keywords: Asset Pricing, Portfolio Efficiency, Conditional Factor Models

1 Introduction

The consumption-based framework for asset pricing, going back to Lucas (1978), is one of the most powerful theoretical paradigms in finance. Most asset pricing models, including the classic CAPM, can be obtained as special cases of, or as proxies for, this model. In addition, the consumption-based framework addresses many of the criticisms leveled at the classic CAPM, such as its failure to account for hedging demands (Merton 1973) or the fact that the market portfolio cannot be proxied by a portfolio of common stocks (Roll 1977). The poor empirical performance of the consumption (C-) CAPM, as documented among others by Hansen and Singleton (1982), and Breeden, Gibbons, and Litzenberger (1989), is thus a puzzle.

In a recent paper, Lettau and Ludvigson (2001b) attempt to resurrect the consumption CAPM by considering a modified version of the model, where the consumption growth factor is scaled by lagged consumption-wealth ratio (CAY), the variable introduced by Lettau and Ludvigson (2001a) and shown to have considerable ability in predicting asset returns[1]. They use the approach of Campbell and Cochrane (2000), which expresses a conditional factor model as an unconditional one in which the factor loadings are constant but the factors themselves are scaled by the conditioning variable. They find that this model out-performs the unscaled versions of the C-CAPM in explaining the cross-section of expected returns on the 25 size and book-to-market portfolios of Fama and French. In particular, they claim that the celebrated 'value premium' can be largely explained by the covariance of an asset's return with scaled consumption growth.

Subsequently a number of studies, for example Hodrick and Zhang (2001), have analyzed the Lettau-Ludvigson framework and found that the model fails various specification error tests, suggesting that the model is mis-specified and thus can have large pricing errors particularly

[1]See also Abhyankar, Basu, and Stremme (2005). Note however that CAY has been criticized on the grounds of the 'look-ahead bias' inherent in its construction, see Brennan and Xia (2005).

out-of-sample, even though it does a reasonable job of explaining the cross-section of expected returns. These studies thus cast doubt as to whether the scaled C-CAPM can indeed be a true asset pricing model. Note that the conditional linear specification, in which asset betas are constant but the factors are scaled by the conditioning instruments, is equivalent to a specification with unscaled factors but time-varying betas. Ghysels (1998) analyzes this latter specification and finds that out-of-sample such models tend to have in fact larger pricing errors than unscaled models.

In this paper, we construct an optimally scaled version of the consumption CAPM and investigate whether it can be a true asset pricing model. We regard the scaled factor model as one in which the factor loadings are time-varying, an approach first advocated by Ferson, Kandel, and Stambaugh (1987), Harvey (1989), and Shanken (1989). We improve upon the existing empirical literature in two ways; first, we do not constrain the factor loadings (betas) to be linear functions of the instruments as advocated in these papers. In fact, our methodology allows us to construct the *optimal*[2] factor loadings as (typically non-linear) functions of the instrument (CAY). Our approach thus gives the model the best possible *ex-ante* chance of success, because the optimal use of the instrument is likely to reduce the specification errors for betas relative to the linear specification. As a consequence, our methodology allows us to assess whether *any* version of the C-CAPM, scaled or not, can ever be a true pricing model. Second, our framework also allows us to assess whether the model prices *actively managed* portfolios correctly, where the portfolio weights are optimal (in the sense of mean-variance efficiency) functions of the instrument. This is important because once the factor risk-premia are allowed to be time-varying functions of some conditioning information, it is unrealistic not to allow the same information to be used in the formation of portfolios. In other words, conditioning information makes pricing models more flexible, but also enlarges the space of assets the model is required to price. The optimal use of conditioning information in portfolio formation was first studied in Hansen and Richard

[2]Here, 'optimal' is defined as minimizing the in-sample pricing errors induced by the scaled model.

(1987), and more recently in Ferson and Siegel (2001).

In order to assess whether a given (set of) factor(s) can give rise to a true asset pricing model, we use a new measure of specification error for scaled factor models, developed in Basu and Stremme (2005). This test exploits the close links between the stochastic discount factor framework and mean-variance efficiency. Specifically, the test measures the distance between the efficient frontier spanned by the factors (or factor-mimicking portfolios) and the frontier spanned by the traded assets. We show that a conditional factor model is a true asset pricing model if and only if the two frontiers coincide, i.e. if and only if our distance measure evaluates to zero. We also show that our test is proportional to the difference in maximum squared Sharpe ratios in the spaces of returns generated by managed portfolios of the traded assets and the factor-mimicking portfolios, respectively. As a consequence, we show that the model is a true asset pricing model if and only if it is possible to construct a dynamically managed strategy, using the factor-mimicking portfolios as base assets, that is unconditionally mean-variance efficient relative to the frontier spanned by the traded assets. This enables us to study the performance of the model both in and out-of-sample. To facilitate a direct comparison with the results of Lettau and Ludvigson (2001b), we also analyze how well the optimally scaled model succeeds in explaining the cross-section of expected returns. It should be pointed out however that the latter is only a necessary and not sufficient condition for the model to be a true asset pricing model. This is because unconditional moments are insufficient to assess conditional pricing errors.

We test the model on two different sets of traded assets; the 5×5 portfolios sorted by size and book-to-market ratio, as used in Lettau and Ludvigson (2001b), as well as 30 portfolios sorted by industry. We find that the C-CAPM, optimally scaled by CAY, can explain about 80% of the cross-section of expected returns on the 5×5 size and book-to-market portfolios, thus out-performing the linearly scaled model considered in Lettau and Ludvigson (2001b). Since the unscaled (C-)CAPM is found to explain no more than 10-20% of the cross-section of expected returns, these results seem to confirm the power of CAY as a scaling instrument. In particular, the optimally scaled C-CAPM indeed seems to explain a large portion of the size and value premia documented famously by Fama and French (1992).

However, when the model is tested on the 30 industry portfolios, even the optimally scaled C-CAPM explains only about 10% of the cross-section of expected returns. In this case, the scaled model in fact performs slightly worse than the corresponding unscaled version. This indicates that the model scaled by CAY, while capturing part of the size and value premia, is nonetheless mis-specified, consistent with the findings of Hodrick and Zhang (2001). Moreover, out-of-sample the scaled model does not succeed in explaining any significant portion of the cross-section of expected returns on either of the two sets of base assets considered.

To assess the performance of the C-CAPM as a conditional asset pricing model, we then construct factor-mimicking portfolios in the asset spaces, using the methodology developed in Basu and Stremme (2005)[3]. We then use these to construct the optimal factor loadings as functions of the conditioning instrument, and evaluate our measure of model specification error. We find that the optimal scaling function for the consumption-growth factor is in fact highly non-linear in the instrument, in contrast to the linear specification that is used predominantly in the existing literature. This explains the superior performance of the optimally scaled model in explaining the cross-section of expected returns, as compared to the linearly scaled model of Lettau and Ludvigson (2001b).

The in-sample estimates of our specification error test show that the optimal use of CAY as scaling instrument indeed significantly improves the performance of the model. In the case of the size and book-to-market portfolios, the optimal use of CAY more than doubles the factor Sharpe ratio (from 0.22 to 0.51). In contrast, the optimal use of CAY in portfolio formation widens the frontier spanned by the base assets only marginally (the Sharpe ratio increases from 1.49 to 1.80). The latter is due to the fact that the size and value effect largely dominates the predictive power of CAY. Our results are quite different in the case where

[3]Similar expressions for factor-mimicking portfolios are also derived, using a slightly different approach, in Ferson, Siegel, and Xu (2005).

the base assets are the 30 industry portfolios. While the Sharpe ratios in all cases are much lower than for the size and book-to-market portfolios, the *relative* effect of introducing CAY as scaling instrument is more dramatic[4]. While the slope of the asset frontier almost doubles (the Sharpe ratio increases from 0.90 to 1.57), the factor Sharpe ratio increases dramatically from only 0.04 to 0.53.

However, while the optimal use of CAY clearly improves the performance of the consumption CAPM, the model nonetheless fails our specification test, indicating that pricing errors are still large. In the case of the size and book-to-market portfolios, the factor-mimicking portfolio achieves only about 40% of the fixed-weight asset Sharpe ratio, and less then 30% of the maximum Sharpe ratio of active portfolios. In other words, even when CAY is used optimally, the model is seriously mis-specified, producing considerable pricing errors even when asked to price only static portfolios. While the model performs slightly better in the case of the 30 industry portfolios, the factor-mimicking portfolio still achieves only about one third of the optimal asset Sharpe ratio. It does, however, achieve about 60% of the fixed-weight Sharpe ratio, indicating that the model comes considerably closer to being able to price the static industry portfolios than it does the size and book-to-market portfolios. In comparison we find that the 'classic' CAPM, while still falling short of being a true asset pricing model, shows considerably better performance. In particular, the factor-mimicking portfolio associated with excess market returns achieves more than 80% of the fixed-weight Sharpe ratio for the 30 industry portfolios. Moreover, while the performance of the C-CAPM deteriorates further out-of-sample, the performance of the classic CAPM is more robust.

Our analysis shows that the optimal use of lagged consumption-wealth ratio does significantly improve the performance of the model. However our conclusion is that while the consumption CAPM scaled by CAY does indeed explain a significant portion of cross-section of expected returns in-sample for the size and book-to-market portfolios, it still has large pricing errors.

[4]These results are consistent with the findings of Abhyankar, Basu, and Stremme (2005), who compare the predictive ability of various conditioning instruments.

Our findings also prove that the linear scaling prevalent in the literature is clearly sub-optimal, leading to larger-than-necessary specification errors. We find that no version of the consumption CAPM scaled by CAY passes our specification test, failing even to price static portfolios. The performance of the model deteriorates considerably out-of-sample which may be due, in part, to the look-ahead bias in the construction of CAY, as observed in Brennan and Xia (2005). Our findings suggest that researchers should look to incorporate additional factors in order to significantly improve the performance of the model. For example, Basu and Stremme (2005) show that the Fama-French 3-factor model, augmented by skewness and kurtosis factors, successfully prices static portfolios.

The remainder of the paper is organized as follows. Section 2, describes the model and establishes our notation. In Section 3, we outline the theoretical methodology and develop our test, while Section 4 focuses on the empirical analysis. Section 5 concludes. The proofs of the mathematical results stated in this paper are available from the authors upon request.

2 Set-Up and Notation

In this section, we define the model and establish our notation. We construct the sets of 'actively managed' portfolios of the base assets and the factor-mimicking portfolios.

2.1 Traded Assets and Managed Pay-Offs

The information flow in the economy is described by a discrete-time filtration $(\mathcal{F}_t)_t$, defined on some probability space $(\Omega, \mathcal{F}, P)$. We fix an arbitrary $t > 0$, and consider the period beginning at time $t-1$ and ending at t. Denote by L_t^2 the space of all $\mathcal{F}_t$-measurable random variables that are square-integrable with respect to P. We interpret Ω as the set of 'states of nature', and L_t^2 as the space of all (not necessarily attainable) state-contingent pay-off claims, realized at time t.

Traded Assets:

There are n traded risky assets, indexed $k = 1 \ldots n$. We denote the gross return (per dollar invested) of the k-th asset by $r_t^k \in L_t^2$, and by $\tilde{R}_t := (\, r_t^1 \, \ldots \, r_t^n \,)'$ the n-vector of risky asset returns. In addition to the risky assets, a risk-free is traded with gross return $r_t^0 = r_f$.

Conditioning Information:

To incorporate conditioning information, we take as given a sub-σ-field $\mathcal{G}_{t-1} \subseteq \mathcal{F}_{t-1}$. We think of $\mathcal{G}_{t-1}$ as summarizing all information on which investors base their portfolio decisions at time $t-1$. In our empirical applications, $\mathcal{G}_{t-1}$ will be chosen as the σ-field generated by lagged consumption-wealth ratio (CAY)[5]. To simplify notation, we write $E_{t-1}(\,\cdot\,)$ for the conditional expectation operator with respect to $\mathcal{G}_{t-1}$.

Managed Portfolios:

We allow for the formation of *managed* portfolios of the base assets. To this end, denote by X_t the space of all elements $x_t \in L_t^2$ that can be written in the form,

$$x_t = \theta_{t-1}^0 r_f + \sum_{k=1}^{n} (\, r_t^k - r_f \,) \theta_{t-1}^k, \tag{1}$$

for $\mathcal{G}_{t-1}$-measurable functions θ_{t-1}^k. To simplify notation, we write (1) in vector form as $x_t = \theta_{t-1}^0 r_f + (\, \tilde{R}_t - r_f e \,)' \, \theta_{t-1}$, where e is an n-vector of 'ones'. We interpret X_t as the space of managed pay-offs, obtained by forming combinations of the base assets with time-varying weights θ_{t-1}^k that are functions of the conditioning information.

Pricing Function:

Because the base assets are defined by their *returns*, we set $\Pi_{t-1}(\, r_t^k \,) = 1$ for $k = 0, 1, \ldots n$, and extend Π_{t-1} to all of X_t by conditional linearity. In particular, for an arbitrary pay-off

[5]Other examples of conditioning variables considered in the literature include, among others, dividend yield (Fama and French 1988), or interest rate spreads (Campbell 1987),

$x_t \in X_t$ of the form (1), it is easy to see that $\Pi_{t-1}(x_t) = \theta^0_{t-1}$. By construction, the pricing rule Π_{t-1} satisfies the 'law of one price', a weak from of no-arbitrage condition.

2.2 Stochastic Discount Factors

We use the stochastic discount factor framework to define what it means for a set of factors to give rise to an admissible asset pricing model.

Definition 2.1 *By an* admissible stochastic discount factor (SDF) *for the model* (X_t, Π_{t-1}), *we mean an element* $m_t \in L^2_t$ *that prices all base assets conditionally correctly, i.e.*

$$E_{t-1}(m_t\, r^k_t) = \Pi_{t-1}(r^k_t) = 1 \qquad \textit{for all } k = 0,\, 1,\, \dots n. \tag{2}$$

The existence of at least one SDF is guaranteed by the Riesz representation theorem, but unless markets are complete it will not be unique. Much of modern asset pricing research focuses on deriving plausible SDFs from principles of economic theory, and then empirically testing such candidates against observed asset returns. Note that in our definition, the SDF is required to price the base assets *conditionally.* The vast majority of asset pricing model tests considered in the literature have used the unconditional version of (2). In our setting, we allow both assets and factors to be dynamically managed, and thus we are testing the conditional version of the pricing equation. Note also that, if m_t is an admissible SDF in the sense of (2), linearity implies $E_{t-1}(m_t\, x_t) = \Pi_{t-1}(x_t)$ for any arbitrary managed pay-off $x_t \in X_t$. In other words, an SDF that prices all base assets correctly is necessarily compatible with the pricing function Π_{t-1} for managed pay-offs. Taking expectations we obtain,

$$E(m_t\, x_t) = E(\Pi_{t-1}(x_t)) =: \Pi_0(x_t). \tag{3}$$

In other words, any SDF that prices the base assets (conditionally) correctly must necessarily also be consistent with the unconditional pricing rule Π_0. In fact, it is easy to show that a candidate m_t is an admissible SDF if and only if (3) holds for all $x_t \in X_t$. We can thus interpret (3) as a set of *moment conditions* that any candidate SDF must satisfy. There are many empirical techniques (e.g. GMM) to estimate and test such restrictions. However, as the space X_t of 'test assets' is infinite-dimensional, such tests will typically yield only necessary but not sufficient conditions for the SDF.

This problem can be overcome by exploiting the close link between the SDF framework and mean-variance efficiency. More specifically, one can obtain necessary *and* sufficient conditions by testing how the candidate SDF acts on the unconditionally efficient frontier in the space X_t of managed pay-offs (see Section 3 below). By two-fund separation, this reduces the test to a one-dimensional problem. Motivated by this observation, we set $R_t = \Pi_0^{-1}\{1\}$. In other words, R_t is the set of all managed pay-offs that have unit price and thus represent the *returns* on dynamically managed portfolios.

2.3 Conditional Factor Models

Our focus here is not the *selection* of factors, but rather the construction and testing of models for a given set of factors. Therefore, we take as given m factors, $F_t^i \in L_t^2$, indexed $i = 1 \ldots m$. Denote by $\tilde{F}_t = (F_t^1, \ldots, F_t^m)'$ the m-vector of factors. In general we do not assume the factors to be traded assets, that is we may have $F_t^i \notin X_t$.

Definition 2.2 *We say that the model* (X_t, Π_{t-1}) *admits a* conditional factor structure, *if and only if there exist* $\mathcal{G}_{t-1}$*-measurable functions* a_{t-1} *and* b_{t-1}^i *such that,*

$$m_t = \alpha_{t-1} + \sum_{i=1}^{m} F_t^i \, b_{t-1}^i \tag{4}$$

is an admissible SDF for the model in the sense of Definition 2.1.

We refer to the coefficients b_{t-1}^i as the *conditional factor loadings* of the model and write (4) in vector notation as $m_t = \alpha_{t-1} + \tilde{F}_t' \, b_{t-1}$. We emphasize that the above specification defines a *conditional* factor model, in that the coefficients a_{t-1} and b_{t-1}^i are allowed to be functions of the conditioning information. In other words, in this specification the conditional risk premia associated with the factors are allowed to be time-varying. This potentially gives the model the flexibility necessary to price also managed portfolios, since the co-efficients of the model can respond to the same information that is used in the formation of portfolios.

Factor-Mimicking Portfolios:

Since the factors need not be traded assets, we construct factor-mimicking portfolios within the space R_t of managed returns.

Definition 2.3 *An element* $f_t^i \in X_t$ *is called a* factor-mimicking portfolio (FMP) *for the factor* $F_t^i \in L_t^2$ *if and only if* $\Pi_{t-1}(\, f_t^i \,) = 1$, *and*

$$\rho^2\big(\, f_t^i,\, F_t^i \,\big) \geq \rho^2\big(\, r_t,\, F_t^i \,\big) \quad \textit{for all } r_t \in X_t \textit{ with } \Pi_{t-1}(\, r_t \,) = 1. \tag{5}$$

Note that we define an FMP via the concept of maximal correlation with the factor. In the literature, it is also common to characterize factor-mimicking portfolios by means of an orthogonal projection[6]. However, it can be shown that these characterizations are in fact equivalent. To define our test, we now take the factor-mimicking portfolios themselves as base assets, and consider the space of pay-offs attainable by forming managed portfolios of FMPs. Specifically, denote by X_t^F the space of all $x_t \in L_t^2$ that can be written in the form,

$$x_t = \phi_{t-1}^0 r_f + \sum_{i=1}^m (\, f_t^i - r_f \,) \phi_{t-1}^i, \tag{6}$$

for $\mathcal{G}_{t-1}$-measurable functions ϕ_{t-1}^i. By construction, $\Pi_{t-1}(\, x_t \,) = \phi_{t-1}^0$ for any $x_t \in X_t^F$ of the form (6). Mimicking the construction in the preceding section, we define the set of *returns* in this space as $R_t^F = R_t \cap X_t^F$.

[6]This is for example the approach taken in Ferson, Siegel, and Xu (2005).

3 Tests of Conditional Factor Models

In this section, we define a new measure of model mis-specification in the presence of conditioning information, as developed in Basu and Stremme (2005). This measure gives rise to a necessary *and* sufficient condition for a given set of factors to constitute a viable asset pricing model. Moreover, we show that our measure is closely related to the shape of the efficient portfolio frontier in the augmented pay-off space.

As a starting point, we take as given an unconditionally efficient benchmark return $r_t^* \in R_t$. Although the results outlined below can be shown to be robust with respect to the choice of benchmark return, we follow Hansen and Jagannathan (1997) and take r_t^* as the return with minimum unconditional second moment in R_t.

Definition 3.1 *For given factors $\tilde{F}_t$, the* model misspecification error *is defined as,*

$$\delta_F := \inf_{r_t \in R_t^F} \sigma^2(\, r_t^* - r_t \,), \tag{7}$$

where R_t^F is the space of managed portfolios of FMPs as defined in (6).

In other words, δ_F measures the minimum variance distance between the efficient benchmark return r_t^* and the return space R_t^F spanned by the factor-mimicking portfolios. The following properties, proven in Basu and Stremme (2005), motivate the interpretation of δ_F as a measure of model mis-specification;

(i) One can show that for given set of factors $\tilde{F}_t$, the model admits a factor structure in the sense of Definition 2.2 if and only if $\delta_F = 0$. In other words, our measure defines a necessary *and* sufficient condition for for conditional factor models.

(ii) By construction, r_t^* attains the maximum Sharpe ratio λ_* in the space R_t of managed returns. One can show that any $r_t^F \in R_t^F$ that attains the minimum in (7) also attains the maximum Sharpe ratio λ_F in the return space R_t^F spanned by the FMPs.

(iii) Moreover, it can be shown that δ_F is proportional to the difference in squared Sharpe ratios, $\lambda_*^2 - \lambda_F^2$. In other words, δ_F measures the distance between the efficient frontiers spanned by the base assets and by the factors, respectively.

As a consequence of (i) and (ii), it follows that a given factor model is a true asset pricing model if and only if it is possible to construct a dynamic portfolio of the FMPs that is unconditionally mean-variance efficient in the asset return space. Thus, our condition is an extension of the Gibbons, Ross, and Shanken (1989) test to the case with conditioning information. In fact, the resulting test statistic is similar to a standard Wald test.

3.1 Factor-Mimicking Portfolios

We now give an explicit characterization of the factor-mimicking portfolios as 'managed' portfolios of the base assets. We define the conditional moments,

$$\mu_{t-1} = E_{t-1}(\, \tilde{R}_t - r_f e\,), \quad \text{and} \quad \Lambda_{t-1} = E_{t-1}\big(\,(\, \tilde{R}_t - r_f e\,)(\, \tilde{R}_t - r_f e\,)'\,\big) \tag{8}$$

In other words, excess returns can be written as $\tilde{R}_t - r_f e = \mu_{t-1} + \varepsilon_t$, where ε_t has zero mean and variance-covariance matrix $\Sigma_{t-1} = \Lambda_{t-1} - \mu_{t-1}\mu_{t-1}'$. Similarly, we denote the mixed conditional moments of the factors by

$$\nu_{t-1} = E_{t-1}\big(\, \tilde{F}_t\,\big), \quad \text{and} \quad Q_{t-1} = E_{t-1}\big(\,(\, \tilde{R}_t - r_f e\,)\tilde{F}_t'\,\big) \tag{9}$$

Note that, if an admissible SDF of the form (4) exists, this implies,

$$0 \equiv E_{t-1}\big(\,(\, \tilde{R}_t - r_f e\,)m_t\,\big) = a_{t-1}\mu_{t-1} + Q_{t-1}b_{t-1}.$$

Conversely, if a_{t-1} and b_{t-1} exist so that $a_{t-1}\mu_{t-1} + Q_{t-1}b_{t-1} = 0$, then m_t in (4) prices all *excess* returns correctly and can hence be modified to be an admissible SDF. In other words, the model admits a conditional factor structure if and only if the image of the conditional linear operator Q_{t-1} contains μ_{t-1}. Basu and Stremme (2005) now show that, for a given factor F_t^i, the corresponding factor-mimicking portfolio can be written as,

$$f_t^i = r_f + \big(\, \tilde{R}_t - r_f e\,\big)'\theta_{t-1}^i \qquad \text{with} \quad \theta_{t-1}^i = \Lambda_{t-1}^{-1}\big(\, q_{t-1}^i - \kappa_i\mu_{t-1}\,\big) \tag{10}$$

where q^i_{t-1} is the column of Q_{t-1} corresponding to factor i, and κ_i is a constant.

3.2 Maximum Sharpe Ratios

In this section, we give explicit expressions for the maximum Sharpe ratios, in the spaces of augmented pay-offs spanned by the base assets and the factors, respectively. Denote by λ_* the maximum Sharpe ratio in the asset return space R_t,

$$\lambda_* = \sup_{r_t \in R_t} \frac{E(\, r_t\,) - r_f}{\sigma(\, r_t\,)}. \tag{11}$$

Similarly, denote by λ_F the corresponding maximum Sharpe ratio in the space R^F_t of managed returns spanned by the factors. Abhyankar, Basu, and Stremme (2005) show that λ_* can be written as $\lambda^2_* = E(\, H^2_{t-1}\,)$, where $H^2_{t-1} = \mu'_{t-1}\, \Sigma^{-1}_{t-1}\, \mu_{t-1}$. This expression extends Equation (16) of Jagannathan (1996) to the case with conditioning information. Similarly, Basu and Stremme (2005) show that the maximum Sharpe ratio λ_F in the return space spanned by the factor-mimicking portfolios can be written as $\lambda^2_F = E(\, H^2_{F,t-1}\,)$, where

$$H^2_{F,t-1} = \mu'_{t-1}\Lambda^{-1}_{t-1}Y_{t-1}\left[\, Y'_{t-1}\Lambda^{-1}_{t-1}\Sigma_{t-1}\Lambda^{-1}_{t-1}Y_{t-1}\,\right]^{-1}Y'_{t-1}\Lambda^{-1}_{t-1}\mu_{t-1}, \tag{12}$$

and $Y_{t-1} = Q_{t-1} - \mu_{t-1}\kappa'$. The main result in Basu and Stremme (2005) states that a given set of factors constitutes a true asset pricing model if and only if $\lambda_F = \lambda_*$. To relate this result to the measure δ_F of specification error, they show that,

$$\delta_F = \left(\frac{r_f}{1 + \lambda^2_*} \right)^2 \cdot (\, \lambda^2_* - \lambda^2_F\,), \tag{13}$$

Thus, δ_F indeed measures distance between the efficient frontiers spanned by managed portfolios of the base assets and the factor-mimicking portfolios, respectively. Since by construction $R^F_t \subseteq R_t$, we always have $\lambda_F \leq \lambda_*$ (and hence $\delta_F \geq 0$), with equality if and only if there exists a portfolio in R^F_t that is unconditionally efficient in the space R_t. In other words, a given factor model is a true asset pricing model if and only if it is possible to construct a managed portfolio from the factor-mimicking portfolios that is efficient in the return space spanned by the traded assets[7].

Finally, Basu and Stremme (2005) explicitly derive the weights of the portfolio that attains the maximum Sharpe ratio λ_F in the factor space and show that, if the model is indeed a true asset pricing model, these weights are in fact proportional to the factor loadings. Moreover, because the weights are chosen optimally, even if the model fails to satisfy the test (that is, even if $\delta_F > 0$), the corresponding factor loadings yield the best possible[8] model that can be constructed from the given set of factors. Moreover, because δ_F is attained by a pair of managed portfolios the weights of which can be explicitly characterized, it lends itself ideally to out-of-sample tests of model performance.

4 Empirical Analysis

In this section, we describe the empirical methodology and data used, and report the results of our analysis.

4.1 Methodology

We specialize the set-up of the preceding sections to the case of a single instrument. Specifically, let y_{t-1} be the given $\mathcal{F}_{t-1}$-measurable conditioning variable (in this case CAY), and set $\mathcal{G}_{t-1} = \sigma(\, y_{t-1}\,)$. For the estimation, we use the de-meaned variable $y^0_{t-1} = y_{t-1} - E\,(\, y_{t-1}\,)$. To compute the conditional moments, we estimate a multivariate predictive regression for

[7] In this sense, our test is very similar in spirit to the 'spanning' test developed in Gibbons, Ross, and Shanken (1989).

[8] Here, 'best possible' is defined as having minimal pricing error.

the factors and (excess) asset returns of the form,

$$\begin{pmatrix} \tilde{R}_t - r_f e \\ \tilde{F}_t \end{pmatrix} = \begin{pmatrix} \mu_0 \\ \nu_0 \end{pmatrix} + \begin{pmatrix} \beta \\ \gamma \end{pmatrix} \cdot y^0_{t-1} + \begin{pmatrix} \varepsilon_t \\ \eta_t \end{pmatrix} \tag{14}$$

where ε_t and η_t are independent of y^0_{t-1} with $E_{t-1}(\,\varepsilon_t\,) = E_{t-1}(\,\eta_t\,) = 0$. Moreover, we assume that the time series of $\{\,\varepsilon_t, \eta_t\,\}$ is independently and identically distributed (iid)[9]. In the notation of Section 3, we can then calculate the conditional moments as,

$$\mu_{t-1} = \mu_0 + \beta y^0_{t-1} \quad \text{and} \quad \Lambda_{t-1} = (\,\mu_0 + \beta y^0_{t-1}\,)(\,\mu_0 + \beta y^0_{t-1}\,)' + E(\,\varepsilon_t \varepsilon_t'\,),$$

$$\nu_{t-1} = \nu_0 + \gamma y^0_{t-1} \quad \text{and} \quad Q_{t-1} = (\,\mu_0 + \beta y^0_{t-1}\,)(\,\nu_0 + \gamma y^0_{t-1}\,)' + E(\,\varepsilon_t \eta_t'\,).$$

Note that due to the iid assumption, we can use unconditional expectations to compute the moments of the residuals ε_t and η_t. For each set of base assets, we estimate (14) and then construct the factor-mimicking portfolios using (10). Using (12), we can then compute the maximum Sharpe ratios λ_* and λ_F directly from the above conditional moments. Alternatively, we can use the results from Basu and Stremme (2005) to construct the managed portfolios that (theoretically) attain these Sharpe ratios, and estimate the unconditional moments of their returns. The latter method is used in our out-of-sample tests to assess the robustness of our in-sample results.

Finally, to compute the model-implied expected returns on the base assets, we use the fact that if m_t is a true SDF of the form (4), then the results of the preceding section imply

$$\mu_{t-1} = -\frac{1}{a_{t-1}}\, Q_{t-1} b_{t-1}.$$

Taking unconditional expectations in the above expression yields the vector of unconditional expected returns implied by the model. We then regress the realized average returns cross-sectionally on the model-implied returns.

[9]Note however that we do *not* assume that the ε_t and η_t are mutually independent, i.e. we do not assume the residual variance-covariance matrix to be diagonal.

4.2 Data

Constrained by the availability of data (in particular the consumption-wealth ratio), we use quarterly data covering the period from January 1960 to December 2002.

Base Assets

We conduct our empirical analysis using two sets of base assets; the 5×5 portfolios sorted on firm size and book-to-market ratio, and the 30 portfolios sorted on industry sector. Monthly data on both sets are available from Kenneth French's web site[10]. The 5×5 size and book-to-market portfolios are constructed as the intersection of 5 portfolios sorted on firm size (market equity), and 5 portfolios sorted on book-to-market ratio. The portfolios are rebalanced at the end of June each year. The 30 industry portfolios are constructed at the end of June each year using the four-digit SIC codes.

Factor Models

To facilitate the comparison with the results of Lettau and Ludvigson (2001b), we focus mainly on the Consumption CAPM, in which the single factor is the growth in aggregate (log) consumption. For comparison, we also consider the 'classic' CAPM, in which the factor is given by the (excess) returns on the market portfolio. We use the Fama-French benchmark factors (available from Kenneth French's web site) to extract the market factor.

Conditioning Instrument

As conditioning instrument, we chose the consumption-wealth ratio as constructed in Lettau and Ludvigson (2001a). The (updated) quarterly data are available from Sydney Ludvigson's web site[11]. In a wide class of forward-looking models, the consumption-aggregate wealth ratio summarizes agents' expectations of future returns to the market portfolio. Thus the variable

[10] `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/`

[11] `http://www.econ.nyu.edu/user/ludvigsons/`

captures expectations without requiring the researcher to observe information sets directly. A log-linear approximation to a representative investor's intertemporal budget constraint shows that the log consumption-wealth ratio may be expressed in terms of future returns to the market portfolio and future consumption growth. This leads to a co-integrating relation between log consumption and log wealth. The log consumption-aggregate wealth ratio is not observable because human capital is not observable. To overcome this obstacle, Lettau and Ludvigson (2001a) reformulate the bivariate co-integrating relation between log consumption and log wealth as a trivariate co-integrating relation involving three observable variables, namely log consumption, c_t, log nonhuman or asset wealth, a_t, and log labor income, y_t. Finally the log aggregate consumption-wealth ratio cay_t is then given by

$$\mathrm{cay}_t = c_t - 0.2711\, a_t - 0.6185\, y_t. \tag{15}$$

It has been pointed out by Brennan and Xia (2005) that this variable suffers from a 'look-ahead bias' that is introduced by estimating the parameters of the co-integrating regression between consumption, asset wealth, and labor income using the entire sample. They thus argue that the in-sample predictive power of this variable cannot be taken as evidence that consumers are able to take account of expected returns on risky assets in making their consumption decisions.

4.3 Results

For illustration only, Figure 1 shows the time series of the conditioning variable CAY over the sample period, compared with the contemporaneous returns on the market portfolio. Without making any claims of statistical significance, the graph seems to indicate that CAY indeed possesses some predictive power, with many of the extreme observations in the two time series coinciding at the one-period lag[12].

[12] For a more statistically rigorous analysis of return predictability using various lagged instruments see Abhyankar, Basu, and Stremme (2005).

Cross-Section of Expected Returns

Figure 2 shows the realized expected returns on the 5×5 size and book-to-market portfolios, graphed against the corresponding model-implied returns. Panel A shows the results for the unscaled C-CAPM, while Panel B reports the results for the C-CAPM optimally scaled by CAY. The reported R^2 are the coefficients of determination in the regression of realized on model-implied returns. The figure shows that optimal scaling indeed improves the performance of the model: while the unscaled model explains less than 10% of the cross-section of expected returns, the R^2 for the optimally scaled model is about 80%. Our results thus seem to confirm the findings of Lettau and Ludvigson (2001b). Note however that our model out-performs theirs by about 10%, which indicates that the assumption of linear scaling is too restrictive (see below for a discussion of the optimal factor loadings).

However, when we repeat the same exercise using the 30 industry portfolios as base assets, the results are very different. As Figure 3 shows, even the optimally scaled C-CAPM explains only about 10% of the cross-section of expected returns. In this case, the scaled model in fact performs slightly worse than the corresponding unscaled version. This indicates that the model scaled by CAY, while capturing part of the size and value premia, is nonetheless mis-specified. This conclusion is also supported by the poor out-of-sample performance of the model. In our out-of-sample test (results not reported), even the optimally scaled C-CAPM had virtually no explanatory power for the cross-section of expected returns (with an R^2 of less than 10%), even in the case of the size and book-to-market portfolios.

Model Specification Error

To assess the performance of the C-CAPM as a conditional asset pricing model, we evaluate our measure of model specification error. We computed both the *ex-ante* values derived directly from the conditional moments using (12), as well as the values derived from the *ex-post* moments of the portfolios that attain the maximum Sharpe ratios. However, because the in-sample results are very similar in both cases, we report only one set of figures in the tables below.

Table 4 shows the in-sample results for the C-CAPM, for both sets of base assets (Panel A reports the results for the size and book-to-market portfolios, Panel B those for the 30 industry portfolios). For both sets of assets, the characteristics of the factor-mimicking portfolio (Table 4.A) are very similar (with expected returns of 5.8 and 5.7%, respectively, and a volatility of about 0.5% in both cases). Note that the *ex-post* betas between factor and mimicking portfolio are close to one (0.90 and 0.92, respectively). This validates the construction of the mimicking portfolios, as we have chosen the constant κ in (10) such that the *ex-ante* beta equals 1.

Moving on to Table 4.B, note first that in the fixed-weight case (without using CAY as conditioning instrument), the asset Sharpe ratio is considerably higher for the size and book-to-market portfolios (1.49) than for the 30 industry portfolios (0.90). In contrast, the difference is much less pronounced (1.80 and 1.57, respectively) when the assets are optimally managed using CAY. In other words, while the introduction of the conditioning instrument enlarges the frontier of the 30 industry portfolios by more than 50%, it has only marginal effect on the frontier spanned by the size and book-to-market portfolios. This is due to the size and value effects which, in the case of the size and book-to-market portfolios, largely dominate the predictive power of CAY.

It is clear from the table that the optimal use of CAY as scaling instrument indeed significantly improves the performance of the C-CAPM (with the factor Sharpe ratios increasing from 0.22 to 0.51 for the size and book-to-market portfolios, and from 0.04 to 0.53 for the 30 industry portfolios, respectively). Similar to the asset frontier, the increase is much more dramatic for the 30 industry portfolios. However, our results also show that even the dramatic increase in Sharpe ratio comes not even close to 'resurrecting' the model: the factor-mimicking portfolios achieve only 28 and 34%, respectively, of the corresponding optimally managed asset Sharpe ratios. This proves that the model, even when scaled optimally by CAY, is still mis-specified and has rather large pricing errors, in particular on actively managed portfolios. Our results thus confirm the earlier findings of Hansen and Singleton (1982), or Breeden, Gibbons, and Litzenberger (1989), that the Consumption CAPM does a poor job of pricing in particular the size and book-to-market portfolios. However, our findings add to the existing empirical literature in that we give the model the 'best chance' by not restricting the factor loadings to be linear functions of the instrument.

Moreover, comparing the optimally scaled factor Sharpe ratios (0.51 and 0.53, respectively) with the fixed-weight asset Sharpe ratios (1.49 and 0.90, respectively), we can conclude that the scaled model does not even succeed in pricing static portfolios conditionally correctly (with the factor achieving only about 34 and 59%, respectively, of the fixed-weight asset Sharpe ratios). Our results show that it is not possible to conclude that the value premium for example can be explained by an asset's covariance with scaled consumption growth. These findings thus broadly confirm (albeit using a very different methodology) those of Hodrick and Zhang (2001), who use the Hansen and Jagannathan (1997) discount factor distance to examine the specification error of the original Lettau-Ludvigson model and find it to be quite seriously mis-specified. This is further emphasized by the poor out-of-sample performance of the model (see below).

Our findings are further illustrated by Figure 5, which shows the efficient frontiers in the case of the 5×5 size and book-to-market portfolios: while the fixed-weight (dotted line) and optimally managed (solid line) asset frontiers are very close to one-another, the frontier spanned by the factor (dashed line) does not even capture some of the base assets (unsurprisingly, it is the small value portfolios that display the strongest performance).

For comparison, we also estimated the 'classic' CAPM, scaled by CAY, where the single factor is given by the excess returns on the market portfolio. The results are reported in Table 8 and Figure 9. We find that the classic CAPM performs significantly better than the consumption CAPM (the factor-mimicking portfolio achieving about 40 and 46% of the optimally scaled asset Sharpe ratios, respectively). More interestingly, the optimally scaled factor achieves up to 80% of the fixed-weight asset Sharpe ratio, indicating that the classic CAPM comes quite close to pricing at least passive portfolios correctly. Moreover, unlike the C-CAPM, the classic CAPM loses little of its performance out-of-sample (see below).

Optimal Factor Loadings

Our framework allows us not only to test if a given (set of) factor(s) can constitute a viable asset pricing model, but also to construct the optimal factor loadings as functions of the conditioning instrument. Figure 6 shows the optimal loadings for the 'corner' elements (small growth and value, and large growth and value) within the matrix of size and book-to-market portfolios. As the graphs show, the optimal loadings are highly non-linear in the instrument, in particular around its mean (0.723). Similarly, Figure 7 shows the coefficient (denoted b_{t-1} in Definition 2.2) of the consumption-growth factor in the optimally scaled stochastic discount factor (SDF). Again, the optimal coefficient is highly non-linear in the conditioning instrument CAY, in particular around its mean. Note however that the non-linearity is less pronounced in the case where the base assets are the 30 industry portfolios (graphs not shown here).

Out-of-Sample Results

Because we can explicitly construct the portfolios that (theoretically) attain the maximum Sharpe ratios in both the asset and the factor return spaces, we are able to assess the out-of-sample performance of the model. To do this, we estimate the conditional moments in-sample, use the results to construct the corresponding efficient portfolios, and then estimate the unconditional moments of these portfolios out-of-sample. The results (for the 5×5 size and book-to-market portfolios) are reported in Table 10. While the performance of both models increases slightly out-of-sample, the performance of the C-CAPM does not match that of the full-sample estimates (Table 4). The results are quite different for the 'classic' CAPM which maintains (in fact slightly exceeds) its in-sample performance (Table 8) out-of-sample. Note also that the out-of-sample performance of the unscaled models is very different from the in-sample estimates, while the results for the optimally scaled model are in general more robust. This latter result should however be taken with caution, as it may be driven in part by the 'look-ahead' bias inherent in the construction of the instrumental variable CAY (Brennan and Xia 2005).

We also investigated the ability of the models to explain the out-of-sample cross-section of expected returns on the base assets (results not reported). We found that, in contrast to the in-sample results (Figures 2 and 3), out-of-sample neither model was able to explain any significant portion of the cross-section of expected returns, for both sets of assets (with an R^2 of less than 5% in the regression of realized out-of-sample returns on in-sample model-implied returns).

5 Conclusion

The consumption-based framework has been one of the theoretical mainstays of asset pricing. However the poor empirical performance of the consumption CAPM has long been a puzzle. Lettau and Ludvigson (2001b) claim that the consumption CAPM can be 'resurrected' by scaling the factor by CAY, the predictive variable constructed in Lettau and Ludvigson (2001a), and show that it does a good job of explaining the expected returns of the 25 portfolios sorted by size and book-to-market. Subsequent studies (Hodrick and Zhang 2001) have shown that their model leads to large pricing errors and is seriously mis-specified.

We re-examine this issue using the method of optimal scaling of factor models developed in Basu and Stremme (2005), which utilizes the predictive variable *optimally*. We use a new measure of specification error also developed in Basu and Stremme (2005), which allows us to analyze the performance of conditional factor models. We find that, while the optimal use of CAY does dramatically improve the performance of the consumption CAPM, the model is unable to price the 25 portfolios sorted by size and book-to-market or the 30 industry portfolios correctly, and is thus quite far from being a true asset pricing model. Our optimally scaled model does a good job of explaining the cross-section of expected returns in-sample, but its performance deteriorates considerably out-of-sample which may be due in part to the look-ahead bias in CAY (Brennan and Xia 2005).

References

ABHYANKAR, A., D. BASU, and A. STREMME (2005): "The Optimal Use of Asset Return Predictability: An Empirical Analysis," working paper, Warwick Business School.

BASU, D., and A. STREMME (2005): "A Measure of Specification Error for Conditional Factor Models," working paper, Warwick Business School.

BREEDEN, D., M. GIBBONS, and R. LITZENBERGER (1989): "Empirical Tests of the Consumption-Oriented CAPM," *Journal of Finance*, 44(2), 231–262.

BRENNAN, M., and Y. XIA (2005): "tay's as good as cay," *Finance Research Letters*, 2(1), 1–15.

CAMPBELL, J. (1987): "Stock Returns and the Term Structure," *Journal of Financial Economics*, 18, 373–399.

CAMPBELL, J., and J. COCHRANE (2000): "Explaining the Poor Performance of Consumption-Based Asset Pricing Models," *Journal of Finance*, 55, 2863–2878.

FAMA, E., and K. FRENCH (1988): "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3–25.

FAMA, E., and K. FRENCH (1992): "The Cross-Section of Expected Returns," *Journal of Finance*, 47, 427–465.

FERSON, W., S. KANDEL, and R. STAMBAUGH (1987): "Tests of Asset Pricing with Time-Varying Expected Risk Premiums and Market Betas," *Journal of Finance*, 42, 201–220.

FERSON, W., and A. SIEGEL (2001): "The Efficient Use of Conditioning Information in Portfolios," *Journal of Finance*, 56(3), 967–982.

FERSON, W., A. SIEGEL, and T. XU (2005): "Mimicking Portfolios with Conditioning Information," forthcoming, Journal of Financial and Quantitative Analysis.

Ghysels, E. (1998): "On Stable Factor Structures in the Pricing of Risk: Do Time-Varying Betas Help or Hurt?," *Journal of Finance*, 53, 549–573.

Gibbons, M., S. Ross, and J. Shanken (1989): "A Test of Efficiency of a Given Portfolio," *Econometrica*, 57(5), 1121–1152.

Hansen, L., and R. Jagannathan (1997): "Assessing Specification Errors in Stochastic Discount Factor Models," *Journal of Finance*, 52, 557–590.

Hansen, L., and S. Richard (1987): "The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models," *Econometrica*, 55(3), 587–613.

Hansen, L., and K. Singleton (1982): "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 55, 1269–1286.

Harvey, C. (1989): "Time-Varying Conditional Covariances in Tests of Asset Pricing Models," *Journal of Financial Economics*, 24, 289–317.

Hodrick, R., and X. Zhang (2001): "Evaluating the Specification Errors of Asset Pricing Models," *Journal of Financial Economics*, 62, 327–376.

Jagannathan, R. (1996): "Relation between the Slopes of the Conditional and Unconditional Mean-Standard Deviation Frontier of Asset Returns," in *Modern Portfolio Theory and its Applications: Inquiries into Asset Valuation Problems*, ed. by S. Saito et al. Center for Academic Societies, Osaka, Japan.

Lettau, M., and S. Ludvigson (2001a): "Consumption, Aggregate Wealth and Expected Stock Returns," *Journal of Finance*, 56(3), 815–849.

Lettau, M., and S. Ludvigson (2001b): "Resurrecting the (C)CAPM: A Cross-Sectional Test when Risk Premia are Time-Varying," *Journal of Political Economy*, 109(6), 1238–1287.

Lucas, R. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445.

Merton, R. (1973): “An Intertemporal Capital Asset Pricing Model,” *Econometrica*, 41(5), 867–887.

Roll, R. (1977): “A Critique of the Asset Pricing Theorys Tests. Part I: On Past and Potential Testability of the Theory,” *Journal of Financial Economics*, 4, 129–176.

Shanken, J. (1989): “Intertemporal Asset Pricing: An Empirical Investigation,” *Journal of Econometrics*, 45, 99–120.

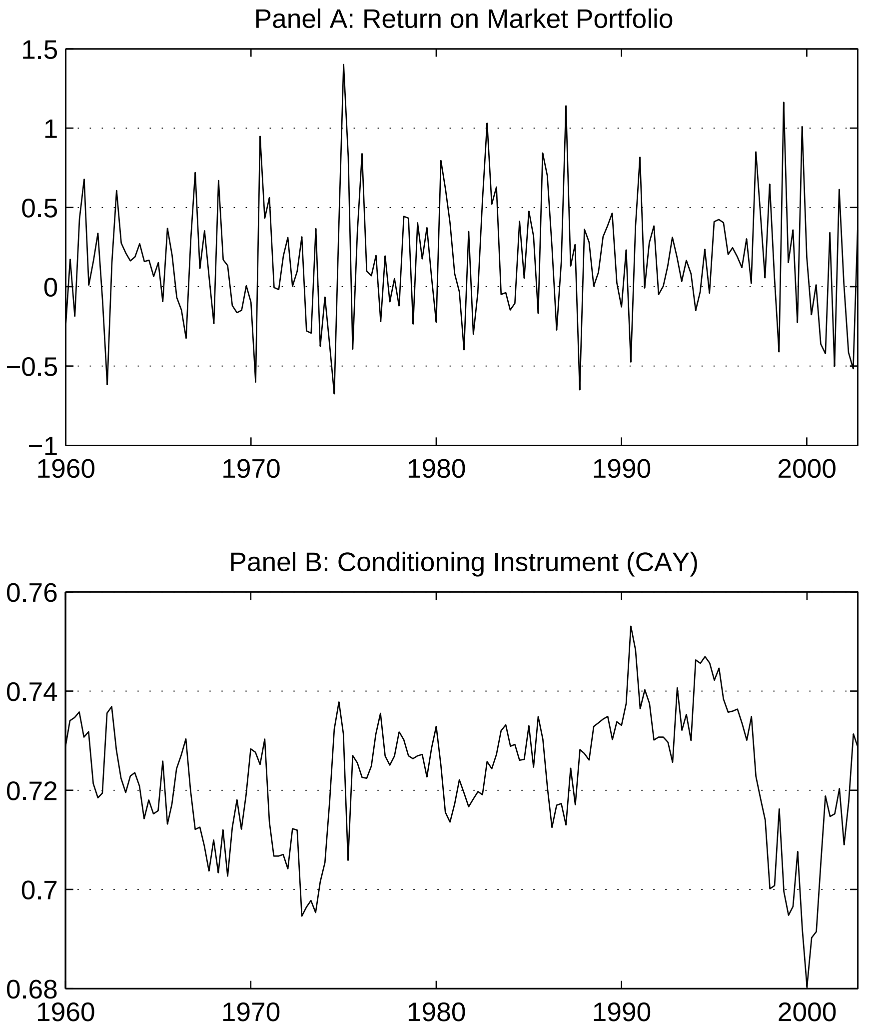

Figure 1: Time Series of Market Returns and CAY

This plot shows the time series of quarterly returns on the market portfolio (Panel A) and the evolution of the consumption-wealth ratio (CAY) that is used as conditioning variable in our empirical analysis. The data for the market return were obtained from Kenneth French's web site.

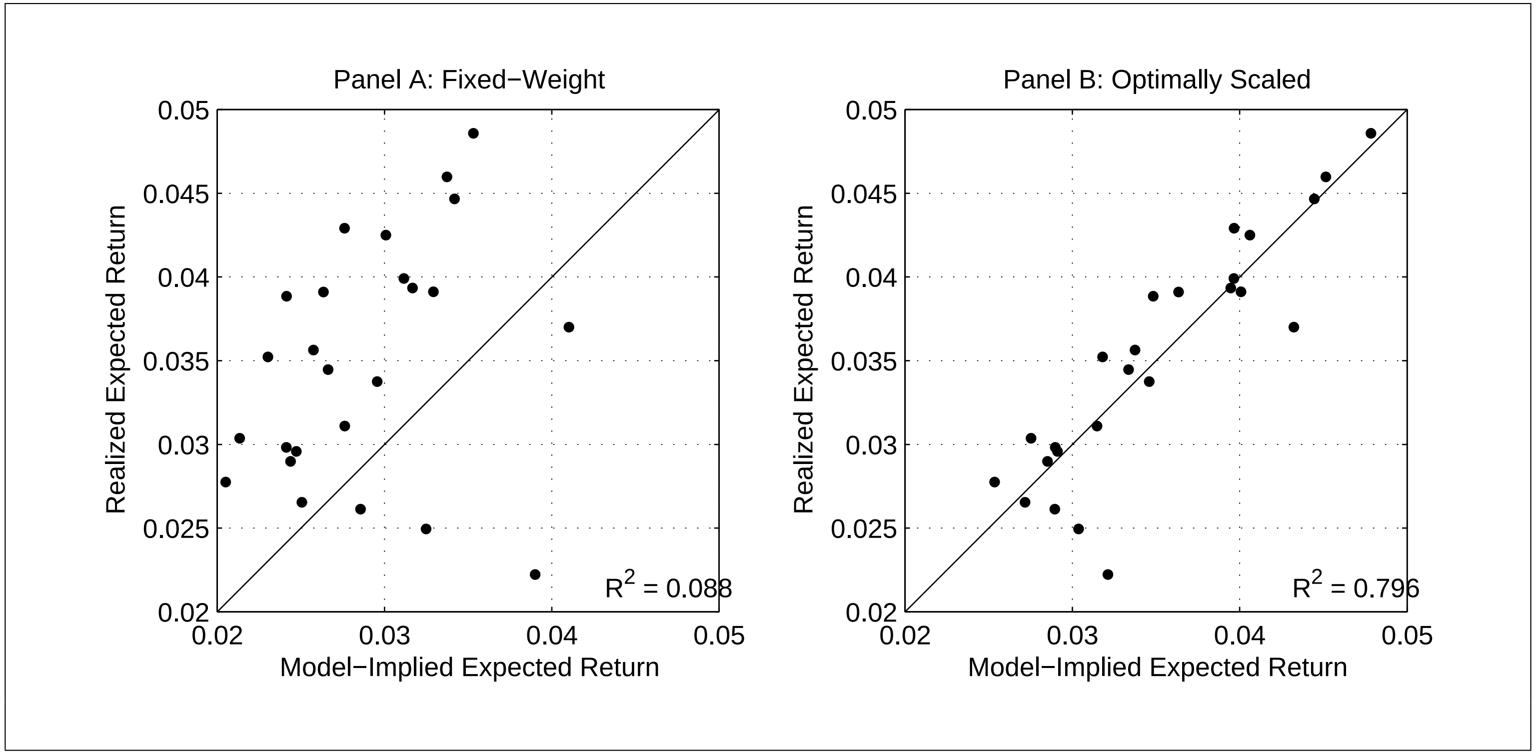

Figure 2: Expected Returns (Consumption CAPM)

This figure graphs the realized average returns (vertical axes) on the base assets against the expected returns implied by the model. The base assets are the 5×5 size and book-to-market portfolios, and the single factor is consumption growth. Panel A shows the returns implied by the unscaled ('fixed-weight') model, while Panel B shows the returns of the model optimally scaled using CAY. Also reported are the R^2 of the cross-sectional regression of realized on model-implied returns.

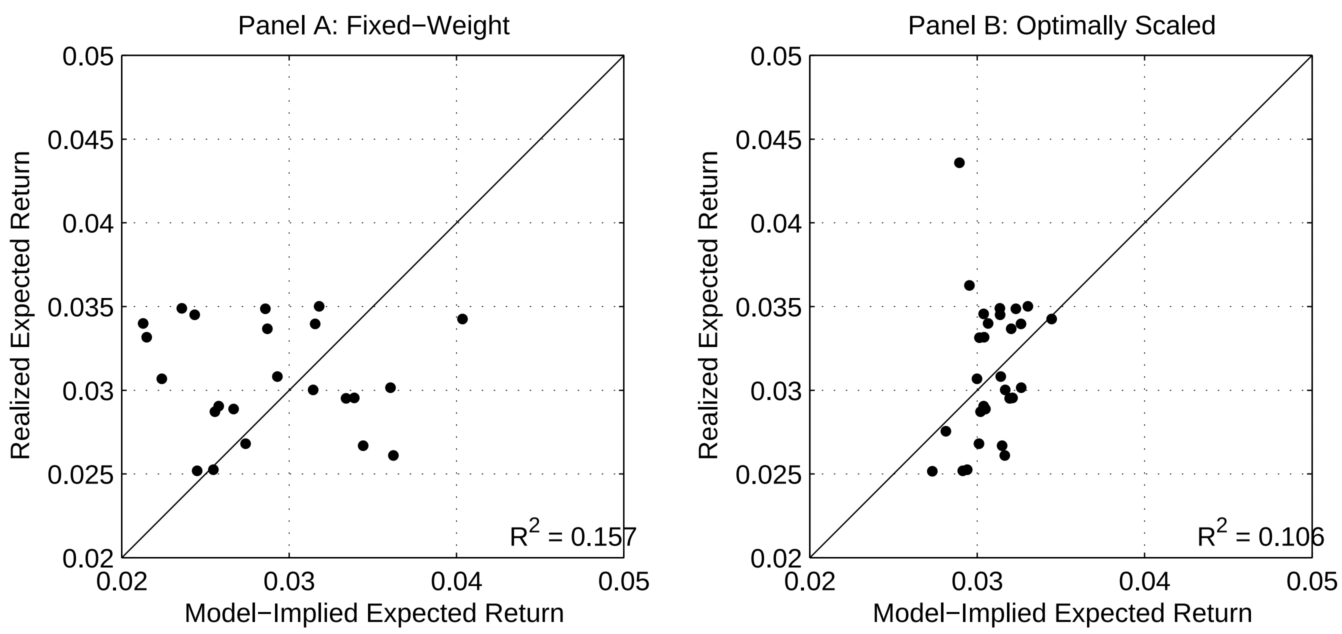

Figure 3: Expected Returns (Consumption CAPM)

This graph is identical to Figure 2, only that the base assets used in this case are the 30 industry portfolios.

	Panel A: FF 25 size/book-to-market			Panel B: FF 30 industry		
	factor-mimicking portf			factor-mimicking portf		
expected return	5.79%			5.70%		
volatility		0.47%			0.48%	
beta (with factor)			0.889			0.922

Table 4.A: Factor-Mimicking Portfolio (Consumption CAPM)

	Panel A: FF 25 size/book-to-market			Panel B: FF 30 industry		
	Sharpe ratio			Sharpe ratio		
	assets	factors	%	assets	factors	%
fixed-weight	1.493	0.223	14.9%	0.901	0.035	3.9%
optimally scaled	1.796	0.509	28.3%	1.574	0.531	33.7%

Table 4.B: In-Sample Estimation Results (Consumption CAPM)

These tables show the in-sample estimation results for the Consumption CAPM. Table 4.A reports the characteristics of the factor-mimicking portfolio associated with the consumption growth factor. Table 4.B shows the maximum Sharpe ratios spanned by the base assets and the factors, respectively, both for the unscaled ('fixed-weight') as well as the optimally scaled model. In each table, Panel A on the left reports the results for the 5 $\times$ 5 size and book-to-market portfolios, while Panel B shows the corresponding results for the 30 industry portfolios. The third column in each panel of Table 4.B shows the fraction of the asset frontier that is spanned by the factors (if this number is 100%, the factor is a true asset pricing model).

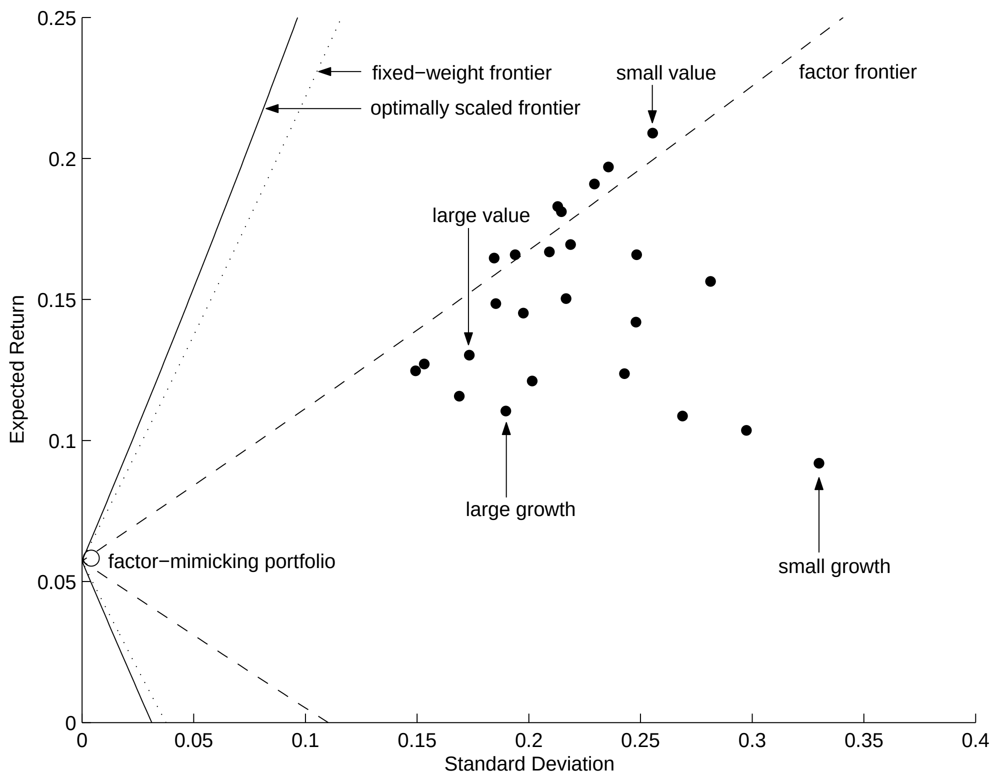

Figure 5: Efficient Frontier (Consumption CAPM)

This plot shows the efficient frontiers generated by the base assets and the factor-mimicking portfolios, respectively. The single factor is consumption growth and the base assets are the 5×5 size and book-to-market portfolios. The base assets (shown as bullets in the figure) are located in the graph from left to right by decreasing size. Within each size group, the different 'styles' are arranged in a 'C'-shaped pattern, with value stocks at the top and growth stocks at the bottom end. The solid and dotted lines show the asset frontier, with and without optimal scaling, respectively. The dashed line shows the frontier spanned by the factors, optimally scaled (the closer the latter comes to the former, the better the factor does at pricing the assets).

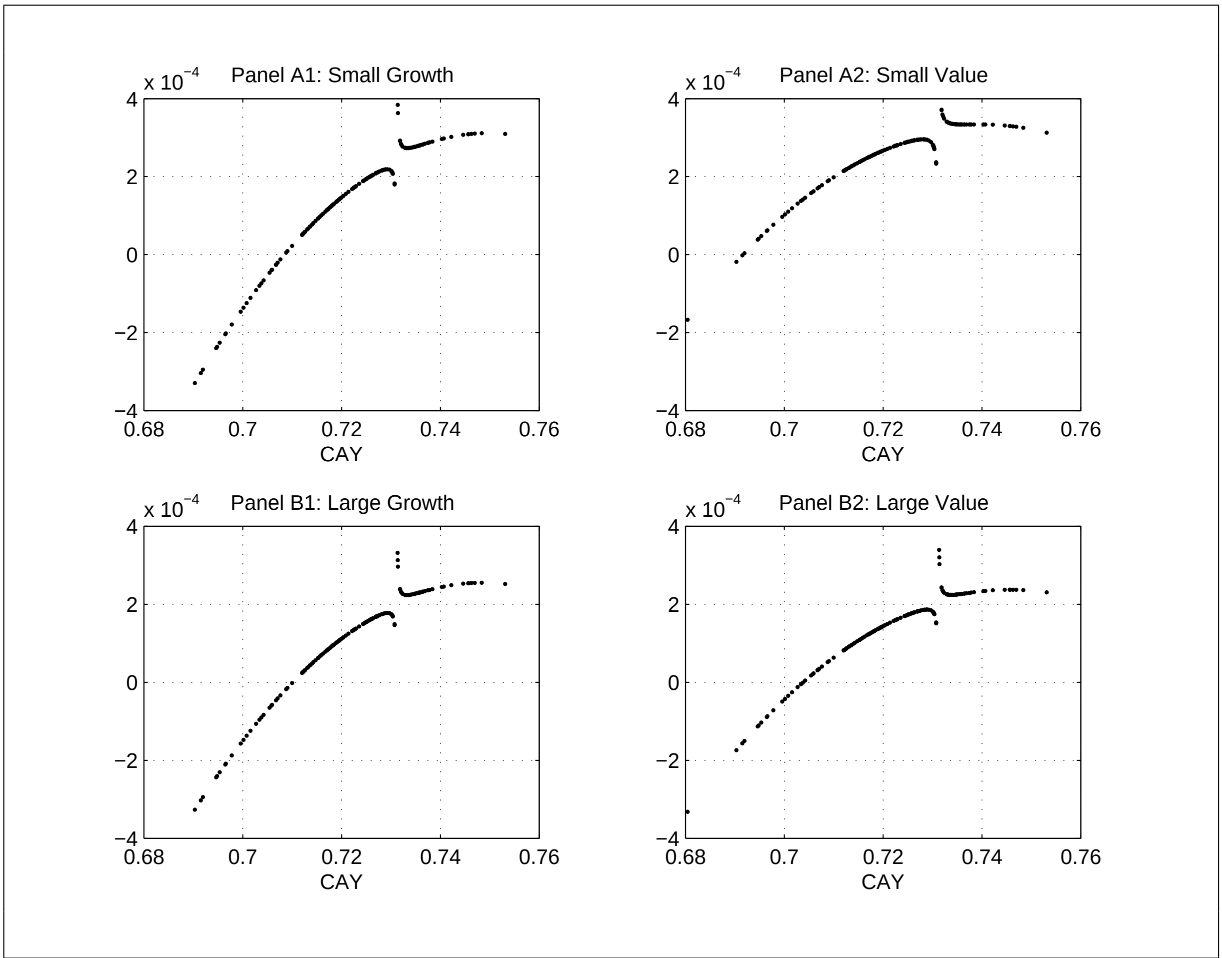

Figure 6: Factor Loadings (Consumption CAPM)

This figure shows the factor loadings (vertical axes) for a selection (the 'corners' of the portfolio matrix, small value, small growth, large value and large growth) of the base assets, as a function of the conditioning instrument CAY (horizontal axes). The base assets are the 5×5 size and book-to-market portfolios, and the single factor is consumption growth.

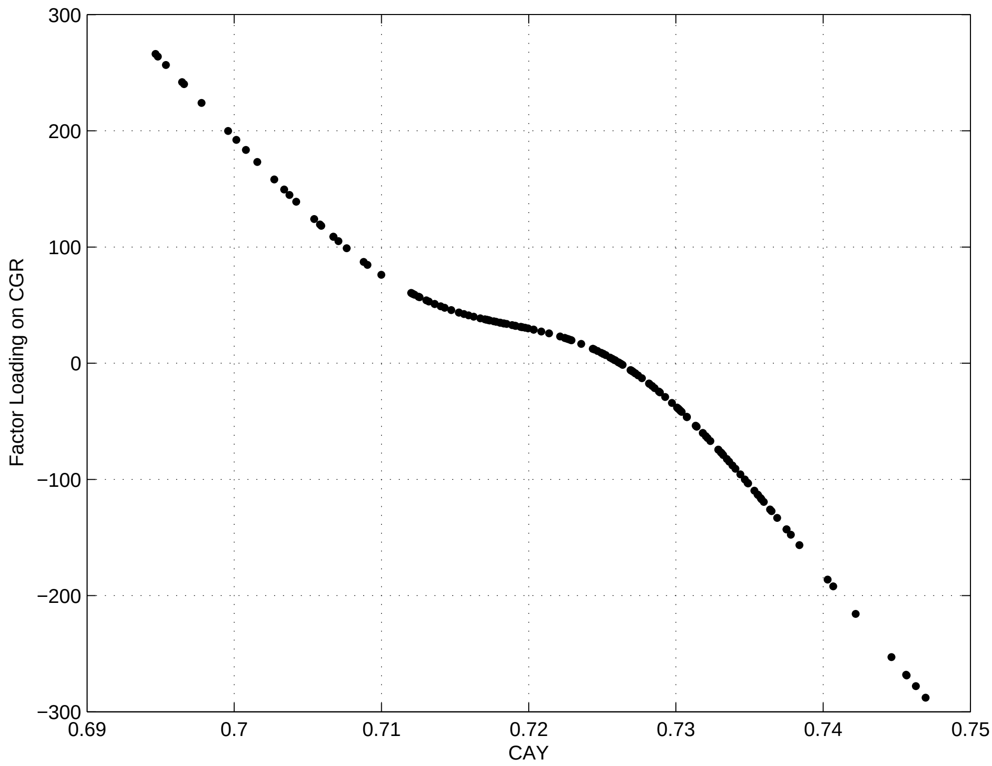

Figure 7: Stochastic Discount Factor (Consumption CAPM)

This figure shows the coefficient (vertical axis) of the single factor that gives the optimal (in the sense of having minimal pricing error) stochastic discount factor, as a function of the conditioning instrument CAY (horizontal axis). The base assets are the 5×5 size and book-to-market portfolios, and the single factor is consumption growth.

	Panel A: FF 25 size/book-to-market	Panel B: FF 30 industry
	factor-mimicking portf	factor-mimicking portf
expected return	12.36%	12.83%
volatility	17.17%	17.17%
beta (with factor)	1.002	1.002

Table 8.A: Factor-Mimicking Portfolio (Classic CAPM)

	Panel A: FF 25 size/book-to-market			Panel B: FF 30 industry		
	Sharpe ratio			Sharpe ratio		
	assets	factors	%	assets	factors	%
fixed-weight	1.493	0.365	24.5%	0.901	0.390	43.3%
optimally scaled	1.796	0.716	39.9%	1.574	0.723	45.9%

Table 8.B: In-Sample Estimation Results (Classic CAPM)

These tables show the in-sample estimation results for the 'classic' CAPM. Table 8.A reports the characteristics of the factor-mimicking portfolio associated with the excess returns on the market portfolio. Table 8.B shows the maximum Sharpe ratios spanned by the base assets and the factors, respectively, both for the unscaled ('fixed-weight') as well as the optimally scaled model. In each table, Panel A on the left reports the results for the 5×5 size and book-to-market portfolios, while Panel B shows the corresponding results for the 30 industry portfolios. The third column in each panel of Table 8.B shows the fraction of the asset frontier that is spanned by the factors (if this number is 100%, the factor is a true asset pricing model).

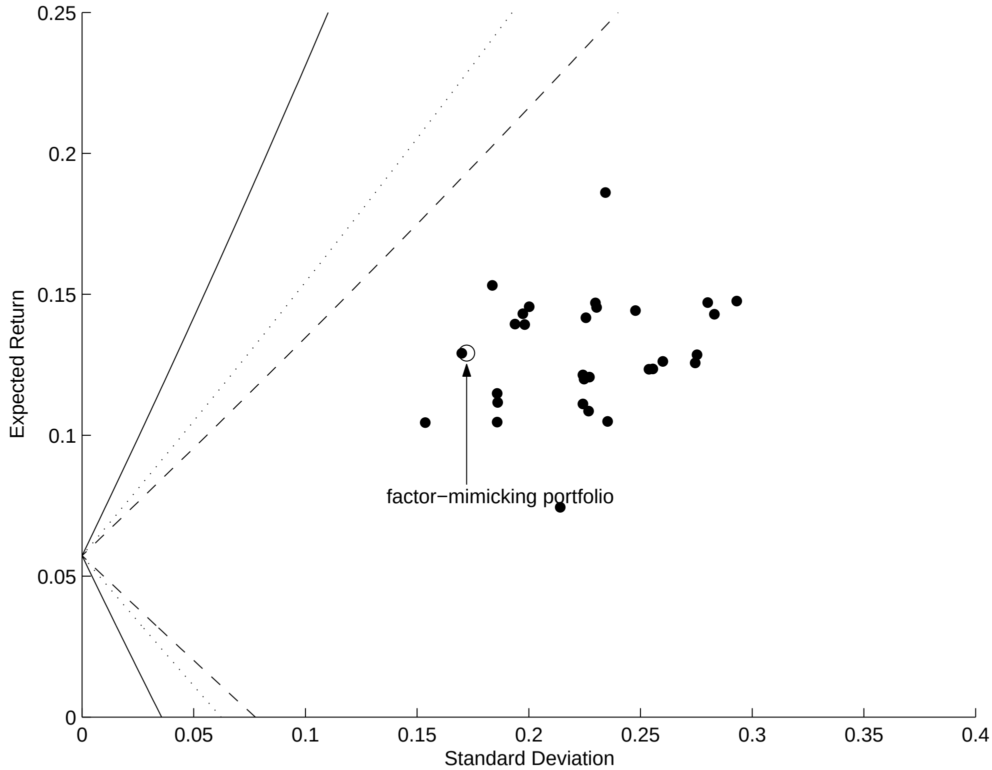

Figure 9: Efficient Frontier (Classic CAPM)

This plot shows the efficient frontiers generated by the base assets and the factor-mimicking portfolios, respectively. The single factor is the excess return on the market portfolio, and the base assets are the 30 industry portfolios. The solid and dotted lines show the asset frontier, with and without optimal scaling, respectively. The dashed line shows the frontier spanned by the factors, optimally scaled (the closer the latter comes to the former, the better the factor does at pricing the assets).

	Panel A: in-sample			Panel B: out-of-sample		
	Sharpe ratio			Sharpe ratio		
	assets	factors	%	assets	factors	%
fixed-weight	1.423	0.002	0.0%	N/A	0.441	31.0%
optimally scaled	1.988	0.327	16.4%	N/A	0.462	23.2%

Table 10.A: Out-of-Sample Estimation Results (Consumption CAPM)

	Panel A: in-sample			Panel B: out-of-sample		
	Sharpe ratio			Sharpe ratio		
	assets	factors	%	assets	factors	%
fixed-weight	1.423	0.289	20.3%	N/A	0.455	32.0%
optimally scaled	1.988	0.713	35.9%	N/A	0.740	37.2%

Table 10.B: Out-of-Sample Estimation Results (Classic CAPM)

These tables show the out-of-sample estimation results for both the Consumption CAPM (Table 10.A) and the 'classic' CAPM (Table 10.B). We estimate the model over the first 25 years of the sample period, and test the resulting model over the remainder of the sample. The base assets are the 5×5 size and book-to-market portfolios, and the conditioning instrument is consumption-wealth ratio (CAY). Each panel shows the maximum Sharpe ratios spanned by the base assets and the factors, respectively, both for the unscaled ('fixed-weight') as well as the optimally scaled model. In each table, Panel A on the left reports the in-sample results, while Panel B shows the corresponding out-of-sample results.