

Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying

Martin Lettau

New York University, Federal Reserve Bank of New York, and Centre for Economic Policy Research

Sydney Ludvigson

New York University

This paper explores the ability of *conditional* versions of the CAPM and the consumption CAPM—jointly the (C)CAPM—to explain the cross section of average stock returns. Central to our approach is the use of the log consumption–wealth ratio as a conditioning variable. We demonstrate that such conditional models perform far better than unconditional specifications and about as well as the Fama-French three-factor model on portfolios sorted by size and book-to-market characteristics. The conditional consumption CAPM can account for the difference in returns between low-book-to-market and high-book-to-market portfolios and exhibits little evidence of residual size or book-to-market effects.

We are grateful to Eugene Fama and Kenneth French for graciously providing the portfolio data; to Nicholas Barberis, John Cochrane, Eugene Fama, Wayne Ferson, Campbell Harvey, John Heaton, Ravi Jagannathan, Timothy Simin, Robert Vishny, and Zhenyu Wang; and to seminar participants at the University of Chicago, the New York Federal Reserve, New York University, the University of Maryland, the University of Rochester, the National Bureau of Economic Research November 1999 Asset Pricing meeting, the American Finance Association 2001 meeting, Utah Winter Finance 2001 conference, and the Western Finance Association July 2000 meeting for many helpful comments. Claire Liou provided able research assistance. The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

[*Journal of Political Economy*, 2001, vol. 109, no. 6]
© 2001 by The University of Chicago. All rights reserved. 0022-3808/2001/10906-0006\$02.50

I. Introduction

Asset pricing theory has fallen on hard times. The capital asset pricing model (CAPM) of Sharpe (1964) and Lintner (1965) has long been a pillar of academic finance, and early evidence seemed to favor the theory's central tenet that the market portfolio be mean-variance efficient (see Black, Jensen, and Scholes 1972; Blume and Friend 1973; Fama and MacBeth 1973). But recently, evidence has mounted that the CAPM is inconsistent with numerous empirical regularities of cross-sectional asset pricing data (see, e.g., Basu 1977; Banz 1981; Shanken 1985; Fama and French 1992, 1993). Perhaps most damning, the CAPM has demonstrated virtually no power to explain the cross section of average returns on assets sorted by size and book-to-market equity ratios (Fama and French 1992, 1993). This failure of hallowed asset pricing theory is displayed in figure 1*a*, which plots the CAPM fitted expected returns for 25 size and book-to-market sorted portfolios against their realized average returns. If the CAPM fit perfectly, all the points in this figure would lie along the 45-degree line. The figure shows clearly that few do.

In response to this failure, Fama and French (1993) develop a three-factor model, with factors related to firm size, book-to-market equity, and the aggregate stock market. They demonstrate, in sharp contrast to the CAPM, its resounding success at capturing the cross section of average returns on these portfolios. This success is reproduced in figure 1*b*. Fama and French (1993, 1995) argue that this three-factor model is successful because it proxies for unobserved common risk in portfolio returns, but this interpretation is controversial since it is not yet clear how these factors relate to the underlying macroeconomic, nondiversifiable risk so proxied.¹

Why has the CAPM failed? One possibility is that its static specification fails to take into account the effects of time-varying investment opportunities in the calculation of an asset's risk. Intertemporal asset pricing models, the most prominent of which is the consumption CAPM (CCAPM), developed by Breeden (1979), initially held out hope of remedying this defect.² Unfortunately, these models have also proved disappointing empirically. The consumption-based model has been rejected on U.S. data in its representative agent formulation with time-separable power utility (Hansen and Singleton 1982, 1983); it has performed no better and often worse than the simple static CAPM in

¹ Liew and Vassalou (2000) provide a first step in establishing an empirical link between the size and book-to-market factors and future macroeconomic variables. They find that, for some countries, these variables have forecasting power for growth in gross domestic product.

² Parts of this theory were also developed by Rubinstein (1976) and Breeden and Litzenberger (1978).

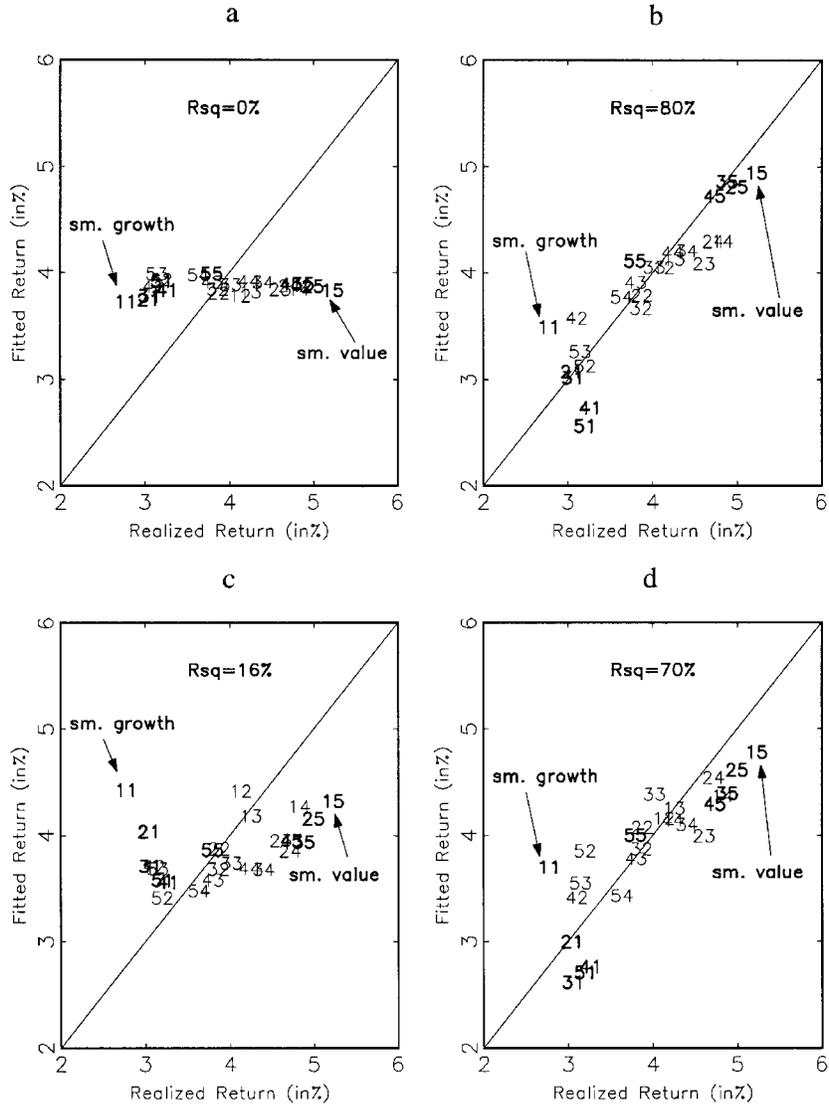


FIG. 1.—Realized vs. fitted returns: 25 Fama-French portfolios: *a*, CAPM; *b*, Fama-French; *c*, consumption CAPM; *d*, consumption CAPM scaled. The figure shows the pricing errors for each of the 25 Fama-French portfolios for the four models. Each two-digit number represents one portfolio. The first digit refers to the size quintiles (1 indicating the smallest firms, 5 the largest), and the second digit refers to book-to-market quintiles (1 indicating the portfolio with the lowest book-to-market ratio, 5 with the highest). The pricing errors are generated using the Fama-MacBeth regressions in table 3 below. The scaling variable is \hat{c}_t .

explaining the cross section of average asset returns (Mankiw and Shapiro 1986; Breeden, Gibbons, and Litzenberger 1989; Campbell 1996; Cochrane 1996); and it has been generally replaced as an explanation for systematic risk by a variety of portfolio-based models (e.g., Elton, Gruber, and Blake 1995; Fama and French 1996). Moreover, the CCAPM performs little better than the CAPM at explaining the cross section of size and book-to-market sorted portfolio returns. Failure number two for asset pricing theory is displayed in figure 1c, which plots the CCAPM fitted returns for the 25 size and book-to-market portfolios against their average realized returns. Again, the contrast with the Fama-French three-factor model is stark.

Despite the empirical shortcomings of the consumption-based model, the reputation of the theoretical paradigm itself remains well preserved. As a measure of systematic risk, an asset's covariance with the marginal utility of consumption has a degree of theoretical purity that is unmatched by other asset pricing models. These other models, including the static CAPM, can almost always be expressed as either special cases of, or proxies for, the consumption-based model.³ Moreover, the consumption-based framework is a simple but powerful tool for addressing the criticisms of Merton (1973), that the static CAPM fails to account for the intertemporal hedging component of asset demand, and Roll (1977), that the market return cannot be adequately proxied by an index of common stocks. According to these rationales, the puzzle is not which model should replace the consumption-based paradigm, but rather why there has been no confirmation of it empirically.

In this paper, we explore a *conditional* version of the consumption CAPM. This conditional model expresses the stochastic discount factor not as an unconditional linear model, as in traditional derivations of the CCAPM, but as a conditional, or *scaled*, factor model. We show empirically that this model explains the cross section of average returns on size and book-to-market sorted portfolios about as well as the Fama-French three-factor model. Figure 1d gives a visual impression of our finding: the scaled CCAPM fitted returns line up along the 45-degree line in a manner that is remarkably similar to the Fama-French three-factor model in figure 1b. (We discuss this figure in more detail below.) Intuitively, conditioning improves the fit of the CCAPM because some stocks are more highly correlated with consumption growth in bad times, when risk or risk aversion is high, than they are in good times, when risk or risk aversion is low. This conditionality on risk premia is missed by unconditional models because they assume that those premia are constant over time.

In this study we also investigate conditional versions of the static CAPM

³ Cochrane (2001) emphasizes this point.

and, following Campbell (1996) and Jagannathan and Wang (1996), of a human capital–augmented CAPM. For each of the models we investigate, precisely the same fundamental factors that price assets in traditional derivations of the CAPM, the human capital CAPM, and the unconditional consumption CAPM are assumed to price assets in this approach. The difference is that factors in the stochastic discount factor are expected to price assets only conditionally, leading to conditional rather than fixed linear factor models.

The results presented here suggest that previous studies documenting the poor empirical performance of both the CAPM and the CCAPM—referred to jointly as the (C)CAPM—may have made inadequate allowances for time variation in the conditional moments of returns. A large and growing body of empirical work finds that expected excess returns on aggregate stock market indexes are predictable, suggesting that risk premia vary over time (see, e.g., Shiller 1984; Campbell and Shiller 1988; Fama and French 1988, 1989; Campbell 1991; Hodrick 1992; Lamont 1998; Lettau and Ludvigson 2001). Yet if risk premia are time-varying, parameters in the stochastic discount factor will depend on investor expectations of future excess returns. For example, in models with habit formation, in which risk premia vary because risk aversion varies, the discount factor will be a state-dependent function of consumption growth.⁴ By contrast, traditional versions of the CAPM and CCAPM imply that these parameters will be state-independent because risk premia are presumed constant.

To capture variation in conditional moments, we explicitly model the dependence of parameters in the discount factor on current-period information, as in Cochrane (1996) and Ferson and Harvey (1999).⁵ This dependence is specified by interacting, or “scaling,” factors with instruments that are likely to be important for summarizing variation in conditional moments. Campbell and Cochrane (2000) argue that such scaled models will perform far better than unscaled models in cross-sectional asset pricing tests if habit formation is present. With this approach, we may express a conditional linear factor model as an unconditional, multifactor model in which the additional factors are simply scaled versions of the original factors. We refer to this version of the (C)CAPM as the *scaled multifactor* model. Scaling factors is one way to incorporate conditioning information; here we use the terms “scaling” and “conditioning” interchangeably.

⁴ Pioneering theoretical work in this area includes Sundaresan (1989), Constantinides (1990), and Campbell and Cochrane (1999).

⁵ This methodology builds off of the work of Ferson, Kandel, and Stambaugh (1987), Harvey (1989), and Shanken (1990), who call for scaling the conditional betas themselves (rather than the factors directly) in a cross-sectional linear regression model in which market betas are expected to vary over time.

The choice of conditioning variable in this study is central to our approach. The linear factor model we consider is a function of investors' conditioning information, which is unobservable. This unobservability is an important practical obstacle to testing conditional factor models since the econometrician's information set is, at best, a subset of the investor's. We argue here, however, that we may largely circumvent this difficulty by using a conditioning variable that summarizes investor expectations of excess returns.

To find such a summary measure of investor expectations, we appeal to a defining feature of any forward-looking model: agents' own behavior reveals much of their expectations about the future. Theories of consumption behavior provide an excellent example. In a wide class of dynamic, optimizing models, log consumption and log aggregate (human and nonhuman) wealth share a common stochastic trend (they are cointegrated), but they may deviate from one another in the short term on the basis of changing expectations of future returns. Accordingly, the log consumption–wealth ratio summarizes investor expectations of discounted future returns to the market portfolio.

The difficulty with this observation is that the consumption–aggregate wealth ratio, specifically the human capital component of it, is not observable. In a recent paper, Lettau and Ludvigson (2001) find that an observable version of this ratio—a cointegrating residual between log consumption, c , log asset (nonhuman) wealth, a , and log labor income, y (referred to subsequently as cay for short)—has striking forecasting power for excess returns on aggregate stock market indexes. In addition, unlike other popular forecasting variables (e.g., the dividend-price ratio), consumption can be thought of as the dividend paid from aggregate wealth, so that movements in the consumption–aggregate wealth ratio summarize expectations about the entire market portfolio, not just the stock market component of it. It follows that cay may have an important advantage over other indicators as a scaling variable in cross-sectional asset pricing tests.

In this paper, we undertake a cross-sectional investigation of the scaled multifactor (C)CAPM using an updated version of the 25 size and book-to-market sorted portfolios constructed in Fama and French (1992, 1993). We focus on these portfolios because explaining their cross-sectional return pattern has presented arguably the greatest empirical challenge to date for theoretically based asset pricing models such as the CAPM and the consumption CAPM.⁶ These models, unlike the Fama-

⁶ The small-firm effect of Banz (1981) notwithstanding, it is well known that the static CAPM is relatively better at explaining the average returns on stock portfolios formed according to size (market value) or industry. In addition, Jagannathan and Wang (1996) show that a conditional CAPM model does a good job of explaining the cross section of returns on portfolios sorted according to size and market betas (covariance with the Center for Research in Security Prices [CRSP] value-weighted index).

French three-factor model, positively fail to explain the strong variation in returns across portfolios in a given size category that differ according to book-to-market equity ratios. If the Fama-French factors truly are mimicking portfolios for underlying sources of macroeconomic risk, there should be some set of macroeconomic factors that performs well in explaining the cross section of average returns on those portfolios. As yet, however, there is little empirical evidence that macroeconomic variables can explain even a small fraction of the variation in these returns.

An important aspect of our results, summarized in figure 1*d*, is that the conditional consumption CAPM, scaled by *cay*, goes a long way toward explaining the celebrated “value premium.”⁷ This suggests that an asset’s risk is determined not by its unconditional correlation with the model’s underlying factor, but rather by its correlation conditional on the state of the economy. The empirical asset pricing literature has become embroiled in controversy over whether this observed value premium is attributable to phenomena captured by firm characteristics (implying a mispricing of value stocks) or to genuine covariance with common risk factors (implying that value stocks are rationally priced).⁸ We show that value stocks earn higher average returns than growth stocks because they are more highly correlated with consumption growth in bad times, when risk premia are high. Thus the results presented here demonstrate that an asset’s covariance with scaled consumption growth can go a long way toward accounting for the value premium, thereby lending support to the view that the reward for holding high-book-to-market stocks arises at least partly as a consequence of true nondiversifiable risk.

The rest of this paper is organized as follows. In Section II we present the general conditional factor model that forms the basis of our empirical work and show how it can be specialized to particular asset pricing models. Next we review the theory in Lettau and Ludvigson (2001) motivating the use of *cay* as a scaling variable. Section III describes the portfolio data and our empirical procedure for testing the (C)CAPM. Section IV presents empirical results on the cross section of average returns. In that section, we compare the performance of conditional factor models in which the return on a value-weighted stock market index or consumption growth is the fundamental factor, with the performance of the simple static CAPM and an unconditional consumption CAPM. We also investigate the cross-sectional explanatory power of a

⁷ The value premium is the well-documented pattern found in average returns that firms with high book-to-market equity ratios have higher average returns than firms with low book-to-market ratios in the same market capitalization category.

⁸ This debate is borne out in several recent papers; see, e.g., Daniel and Titman (1997), Ferson, Sarkissian, and Simin (1999), and Davis, Fama, and French (2000).

scaled human capital CAPM that includes as factors both the return on a value-weighted stock market index and the labor income growth measure advocated by Jagannathan and Wang (1996). The performance of all these models is compared to that of the three-factor model advocated by Fama and French (1993, 1995). We then move on to discuss the average pricing errors for each model. We provide intuition for why the scaled factor models work better than the unscaled models and show that value stocks have higher conditional consumption betas in bad times than growth stocks, suggesting that the former are indeed riskier than the latter. Tests of the conditional factor models when portfolio characteristics such as size and book-to-market ratios are included as additional explanatory variables are also discussed. Finally, in Section IVG, we present results when scaled returns are added to the set of 25 size and book-to-market sorted unscaled returns. Section V discusses alternative estimation methodologies. Section VI presents conclusions.

II. Linear Factor Models with Time-Varying Coefficients

We begin by imposing virtually no theoretical structure, appealing instead to a well-known existence theorem to motivate our empirical approach (see Harrison and Kreps 1979). This theorem states that, in the absence of arbitrage, there exists a stochastic discount factor, or pricing kernel, M_{t+1} , such that, for any traded asset with a net return at time t of $R_{i,t+1}$, the following equation holds:

$$1 = E_t[M_{t+1}(1 + R_{i,t+1})], \quad (1)$$

where E_t denotes the mathematical expectation operator conditional on information available at time t , $M_{t+1} = a_t + b_t R_{e,t+1}$, and $R_{e,t+1}$ is a return on the unobservable mean-variance efficient frontier. We refer to models of the form $M_{t+1} = a_t + b_t R_{e,t+1}$ as *conditional linear factor models*. Special cases of these models with constant coefficients, for example, $M_{t+1} = a + bR_{e,t+1}$, will be referred to hereafter as *unconditional linear factor models*.

It is straightforward to show that the conditional linear factor model given above implies a conditional beta representation given by

$$E_t R_{i,t+1} = R_{0,t} - b_t R_{0,t} \text{Var}_t(R_{e,t+1}) \beta_{ib}$$

where $R_{0,t}$ is the return on a “zero-beta” portfolio uncorrelated with M_{t+1} ,

$$b_t = -\frac{E_t R_{e,t+1} - R_{0,t}}{R_{0,t} \text{Var}_t(R_{e,t+1})},$$

and

$$\beta_{it} = \frac{\text{Cov}_t(R_{e,t+1}, R_{i,t+1})}{\text{Var}_t(R_{e,t+1})}.$$

If conditional moments are time-varying, the parameter b_t in the stochastic discount factor will in general not be constant. Although predictable movements in volatility may be a source of variation in b_t , they appear to be more concentrated in high-frequency data (e.g., Christoffersen and Diebold 2000). Since risk-free interest rates are also not highly variable, the denominator of b_t is not likely to vary much in monthly or quarterly data. On the other hand, a large empirical literature (cited in the Introduction) documents that excess returns are forecastable. Therefore, asset pricing tests that are implemented using monthly or, as in this paper, quarterly data should allow for the possibility of time variation in b_t . In this paper we focus on time variation in equity premia as a source of variation in b_t .

When $M_{t+1} = a_t + b_t R_{e,t+1}$ is plugged into (1) and unconditional expectations are taken, it is straightforward to demonstrate that the conditional model in (1) does not necessarily imply an unconditional version in which a_t and b_t are constants. Following Cochrane (1996), we test conditional factor pricing models of the form given above by explicitly modeling the dependence of the parameters a_t and b_t on a time t information variable, z_t , where z_t is a forecasting variable for excess returns.⁹ (We discuss our choice of conditioning variable further in the next section.) In particular, we may *scale* the factors with instruments containing time t information by modeling the parameters as linear functions of z_t , $a_t = \gamma_0 + \gamma_1 z_t$ and $b_t = \eta_0 + \eta_1 z_t$. Plugging these equations into M_{t+1} above allows us to rewrite a conditional linear factor model as a *scaled multifactor model* with constant coefficients taking the form

$$\begin{aligned} M_{t+1} &= (\gamma_0 + \gamma_1 z_t) + (\eta_0 + \eta_1 z_t) R_{e,t+1} \\ &= \gamma_0 + \gamma_1 z_t + \eta_0 R_{e,t+1} + \eta_1 (z_t R_{e,t+1}). \end{aligned} \quad (2)$$

It follows that the scaled multifactor model can be tested using unconditional moments by rewriting (1) as an unconditional three-factor model with constant coefficients γ_0 , γ_1 , η_0 , and η_1 in the form

$$1 = E[(\gamma_0 + \gamma_1 z_t + \eta_0 R_{e,t+1} + \eta_1 (z_t R_{e,t+1}))(1 + R_{i,t+1})]. \quad (3)$$

⁹ The specification can be easily extended to allow for multiple conditioning variables.

A. *Application of Conditional Factor Pricing to the (C)CAPM*

The derivation above is useful for demonstrating how one can test models in which factors price assets conditionally, but the framework itself contains little theoretical content. In order to test particular theories, we need to place more structure on the discount factor M_{t+1} , and in particular on the choice of reference return, $R_{e,t}$. In the (C)CAPM theories, the true mean-variance efficient reference return may be written as a conditional linear combination of relevant fundamental factors, where the j th factor is denoted f_{jt} . In the Fama and French (1993) specification, a vector of factors, \mathbf{f}_t , contains three portfolio returns. We discuss these models, each a special case of the broader class of scaled multifactor models, in more detail below.

To describe the class of scaled multifactor models more comprehensively, we use vector notation.¹⁰ Denote the vector $\mathbf{F}_{t+1} = (1, z_t, \mathbf{f}'_{t+1}, \mathbf{f}'_{t+1}z_t)'$, or separating out the variable factors z_t , \mathbf{f}_{t+1} , and $\mathbf{f}_{t+1}z_t$ and denoting these together as $\bar{\mathbf{f}}_{t+1}$, write $\mathbf{F}_{t+1} = (1, \bar{\mathbf{f}}'_{t+1})'$. We shall refer to \mathbf{f}_{t+1} as *fundamental factors* (e.g., the market return and consumption growth). The stochastic discount factor of the scaled multifactor representation for each model can be expressed as $M_{t+1} = \mathbf{c}'\mathbf{F}_{t+1}$, where the constant vector $\mathbf{c} = (\gamma_0, \mathbf{b})'$, γ_0 is a scalar, and $\mathbf{b} = (\gamma_1, \eta'_0, \eta'_1)'$ is the vector of constant coefficients on the variable factors, \mathbf{f}_{t+1} . This representation for M_{t+1} implies an unconditional multifactor beta representation for asset i with constant betas given by

$$E[R_{i,t+1}] = E[R_{0,t}] + \boldsymbol{\beta}'\boldsymbol{\lambda}, \quad (4)$$

where $E[R_{0,t}]$ is the average return on a zero-beta portfolio that is uncorrelated with the stochastic discount factor (Black 1972), and $\boldsymbol{\beta} \equiv \text{Cov}(\mathbf{f}, \bar{\mathbf{f}}')^{-1} \text{Cov}(\mathbf{f}, R_{i,t+1})$ is a vector of regression coefficients from a multiple regression of returns on the variable factors. In the empirical analysis that follows, we focus on this Black version of the (C)CAPM (which does not assume the existence of a risk-free security) and freely estimate the constant $E[R_{0,t}]$ as part of the cross-sectional model.

Given (4), it is straightforward to show that

$$\boldsymbol{\lambda} = -E[R_{0,t}] \text{Cov}(\bar{\mathbf{f}}, \bar{\mathbf{f}}')\mathbf{b}. \quad (5)$$

It is important to note that the individual λ_j coefficients in (5) from the scaled multifactor versions of the (C)CAPM do not have a straightforward interpretation as a risk price. To see why, notice that, for each scaled multifactor model, there is an associated conditional factor model from which the scaled multifactor model is derived. For example, the conditional CAPM factor model would be specified $M_{t+1} = a_t +$

¹⁰ This discussion follows the derivation in Cochrane (1996).

$b_t R_{vw,t+1}$ (where $R_{vw,t+1}$ is a proxy for the market return), from which we derive the scaled multifactor model, $M_{t+1} = \gamma_0 + \gamma_1 z_t + \eta_0 R_{vw,t+1} + \eta_1 (z_t R_{vw,t+1})$, using the conditioning information, z_t . More generally, given a conditional linear factor model of the form $M_{t+1} = \mathbf{c}_t'(1, \mathbf{f}'_{t+1})'$, where $\mathbf{c}_t = (a_p, \mathbf{b}'_t)'$, the conditional beta representation for this model is given by analogy to (4) as $E_t[R_{i,t+1}] = R_{0,t} + \tilde{\boldsymbol{\beta}}'_t \tilde{\boldsymbol{\lambda}}_p$ where $R_{0,t}$ is again the return on a zero-beta portfolio, $\tilde{\boldsymbol{\beta}}_{t+1} \equiv \text{Cov}_t(\mathbf{f}_{t+1}, \mathbf{f}'_{t+1})^{-1} \text{Cov}_t(\mathbf{f}_{t+1}, R_{i,t+1})$, and $\tilde{\boldsymbol{\lambda}}_t$ is the vector of period t risk prices of the fundamental factors

$$\tilde{\boldsymbol{\lambda}}_t = -R_{0,t} \text{Cov}_t(\mathbf{f}_{t+1}, \mathbf{f}'_{t+1}) \mathbf{b}_t \quad (6)$$

The period t risk prices in (6) bear no simple relation to the coefficients $\boldsymbol{\lambda}$. Moreover, the methodology employed here (and discussed in more detail later) does not produce estimates of $\tilde{\boldsymbol{\lambda}}_t$. Instead, we estimate cross-sectional regressions of the form (4), which delivers estimates of $\boldsymbol{\lambda}$. Estimates of $\boldsymbol{\lambda}$ from (5) can be used to uncover $\mathbf{b}' = -\boldsymbol{\lambda}' [E[R_{0,t}] \text{Cov}(\bar{\mathbf{f}}, \bar{\mathbf{f}}')]^{-1}$, which we may combine with the definition of $\mathbf{b}_t = \boldsymbol{\eta}_0 + \boldsymbol{\eta}_1 z_t$ to obtain an estimate of \mathbf{b}_t . Without making further assumptions, however, we cannot compute the risk prices for the fundamental factors, $\tilde{\boldsymbol{\lambda}}_p$ because we do not estimate the conditional covariance, $\text{Cov}_t(\mathbf{f}_{t+1}, \mathbf{f}'_{t+1})$, in (6).

Equation (6) shows that the value of $\tilde{\boldsymbol{\lambda}}_p$, the vector of risk prices for each fundamental factor in \mathbf{f}_p depends on \mathbf{b}_t . Our linear specification $\mathbf{b}_t = \boldsymbol{\eta}_0 + \boldsymbol{\eta}_1 z_t$ presumes that fluctuations in \mathbf{b}_t are primarily driven by fluctuations in risk premia and implies a linear forecasting equation for excess returns. While evidence in Lettau and Ludvigson (2001) suggests that these forecasting equations do a good job of picking up fluctuations in future excess returns, as with any linear forecasting model, there are periods in which the model predicts a negative excess return. Since $\tilde{\boldsymbol{\lambda}}_t$ inherits the properties of these linear forecasting models, the value of $\tilde{\boldsymbol{\lambda}}_t$ may change sign from time to time. This aspect of the prediction equation is purely a result of the linear regression specification and is not unique to the use of any particular forecasting variable, z_t . The linear forecasting equations we use below do predict a positive risk premium on average, however, so that it is reasonable to expect that the *average* risk price on each fundamental factor in the conditional model be nonnegative, that is, $E[\tilde{\boldsymbol{\lambda}}_t] \geq 0$.

However, this condition does *not* imply that the individual λ_j coefficients from the scaled multifactor representation in (5) should be nonnegative. In the case of a single-factor model such as the static CAPM, the average risk price for the market beta will have the opposite sign of the average value of \mathbf{b}_t (see [6]) and will be positive as long as the average risk premium is positive. In this case, the conditional covariance

term in (6) will simply be a conditional variance for the value-weighted return. For models with multiple factors, the conditional covariance is not simply a conditional variance, and the average price of risk need not have the opposite sign of the average value of \mathbf{b}_i if the factors are not orthogonal. Nevertheless, if we assume that the conditional covariances and the average zero-beta rate in (6) are approximately constant and we use the specification $\mathbf{b}_i = \boldsymbol{\eta}_0 + \boldsymbol{\eta}_1 z_p$, we may compute a value for the average risk price, $E[\tilde{\lambda}_i]$, of each fundamental factor in the associated conditional factor model using (6). For each of the models investigated below, we make this assumption and check whether, conditional on it, the estimated values of λ imply that $E[\tilde{\lambda}_i] \geq 0$.

We now move on to discuss the special cases of (4) that correspond to the particular scaled multifactor asset pricing models of the CAPM and the CCAPM.

1. The Consumption CAPM

Consider a representative agent economy in which all wealth, including human wealth, is tradable. Let W_t be aggregate wealth (human and nonhuman) in period t , C_t consumption, and $R_{m,t+1}$ the net return on aggregate wealth, or the market portfolio. Subject to an accumulation equation for aggregate wealth, investors maximize the present discounted value of instantaneous utility functions, $u(C_t, X_t)$, where C_t is consumption and X_t captures other factors (e.g., a habit level) that may influence an investor's utility. The first-order conditions for optimal consumption choice are simply special cases of (1), where the equation holds for every asset in the market portfolio and the discount factor, $M_{t+1} \equiv \delta[u_t(C_{t+1}, X_{t+1})/u_t(C_t, X_t)]$, is the intertemporal marginal rate of substitution, with δ the subjective rate of time preference.

Instead of specifying a particular functional form for marginal utility, we assume that M may be approximated as a linear function of consumption growth:

$$M_{t+1} \approx a_t + b_t \Delta c_{t+1}, \quad (7)$$

where a_t and b_t are (potentially time-varying) parameters and Δc_{t+1} is consumption growth, the single fundamental factor in the asset pricing model. Throughout this paper, we use lowercase letters to denote logarithms of variables written in uppercase; for example, $c_t = \ln C_t$. In the notation above, this specification of the CCAPM has a single factor, $f_t = \Delta c_t$, and time variation in the coefficients is modeled by interacting consumption growth with an instrument z_t so that $\mathbf{F}_{t+1} = (\mathbf{1}, z_p, \Delta c_{t+1}, \Delta c_{t+1} z_t)'$.

Regardless of the particular functional form of the investor's utility function, the discount factor can always be expressed as an approximate

linear function of consumption growth by taking a first-order Taylor expansion of M . Examples include time-separable power utility with constant relative risk aversion, $u(C_t) = C_t^{1-\gamma}/(1-\gamma)$, in which case $M_{t+1} \approx \delta(1 - \gamma\Delta c_{t+1})$, and the coefficients a_t and b_t in (7) are constant; or the habit persistence framework of Campbell and Cochrane (1999), $u(C_t, X_t) = (C_t - X_t)^{1-\gamma}/(1-\gamma)$, in which case M_{t+1} takes the form

$$M_{t+1} \approx \delta\{1 - \gamma g \lambda(s_t) - \gamma(\phi - 1)(s_t - \bar{s}) - \gamma[1 + \lambda(s_t)]\Delta c_{t+1}\}, \quad (8)$$

where X_t is the external consumption habit; s_t is the log of the surplus consumption ratio, defined as $S_t \equiv (C_t - X_t)/C_t$; γ is a parameter of utility curvature; g is the mean rate of consumption growth; ϕ is the persistence of the habit stock; and $\lambda(s_t)$ is the sensitivity function specified in Campbell and Cochrane. Similar but more complicated expressions can be derived for internal habit formation models (e.g., Sundaresan 1989; Constantinides 1990) by taking a local linear approximation of the respective stochastic discount factors.

To our knowledge, almost all the cross-sectional tests of the CCAPM to date have pertained to models that assume that the parameters a_t and b_t in (7) are constant.¹¹ Thus they implicitly assume that risk premia are constant, a presumption that produces an asset pricing model in which consumption growth, Δc_{t+1} , is the single factor. But a model like that in equation (8), in which the parameters a_t and b_t are not constant, will have factors in addition to consumption growth when expressed as a scaled factor model. Both the coefficient that multiplies consumption growth and the “intercept” in (8) vary over time and are a function of the surplus consumption ratio, which governs risk aversion. Although these coefficients may be a function of unobservable variables, such as $\lambda(s_t)$ in (8), their fluctuations should be well captured by suitable proxies for time-varying risk premia. In contrast to traditional, discrete-time derivations of the conditional CAPM (e.g., Jagannathan and Wang 1996), which produces just two betas (one for the fundamental factor and one for the risk premium), habit models imply the existence of at least one beta for the multiplicative “cross term” of the factor (consumption growth) and the risk premium.

2. The CAPM and the Human Capital CAPM

A standard derivation of the static CAPM would require simply replacing Δc_{t+1} in (7) with the return to the market portfolio as the relevant factor,

¹¹ An exception is the model of Ferson and Harvey (1993), who estimate a consumption-based asset pricing model using a small cross section of returns including a Treasury bill rate, a government bond rate, a corporate bond rate, a value-weighted stock return, and the return on an index of small stocks. They use lags of consumption growth and a real Treasury bill rate as conditioning variables.

$f_{t+1} = R_{m,t+1}$. The market portfolio is typically proxied by the return on an index of common stocks, but this practice has been challenged by Roll (1977), who argues that such proxies ignore the human capital component of aggregate wealth. Following Mayers (1972) and Fama and Schwert (1977), Campbell (1996) and Jagannathan and Wang (1996) argue that labor income growth may proxy for the return to human capital and find that it has a statistically significant risk price in cross-sectional tests of the CAPM. This specification of the CAPM, explicitly accommodating human capital, would have two factors, the return on a value-weighted stock index, $R_{vw,t}$ and labor income growth, Δy_t , implying that $M_{t+1} = a_t + b_{vw}R_{vw,t+1} + b_{\Delta y}\Delta y_{t+1}$. We refer to this specification as the *human capital CAPM*. To model possible time variation of the parameters in M_{t+1} , a scaled multifactor model given by $\mathbf{F}_{t+1} = (1, z_t, R_{vw,t+1}, \Delta y_{t+1}, R_{vw,t+1}z_t, \Delta y_{t+1}z_t)'$ can be specified in analogy to the consumption model outlined above.

A critical consideration in using the scaled multifactor approach to test the (C)CAPM, or any conditional asset pricing model, is the choice of conditioning variable, z_t . The discussion above suggests that the parameters a_t and b_t will depend on risk premia. Thus we seek a scaling variable that provides a summary measure of expected excess returns. The next subsection describes our choice of conditioning variable, a proxy for the log consumption–wealth ratio, and discusses how it furnishes such a summary by providing a brief overview of the results in Lettau and Ludvigson (2001).

B. The Conditioning Variable

Why use the consumption–wealth ratio as a conditioning variable? First, 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 captures expectations without requiring the researcher to observe information sets directly. Second, as Cochrane (2001, chap. 8) emphasizes, the CAPM can be derived from several special cases of the CCAPM. These special cases also often imply that the parameters in the stochastic discount factor of the CAPM (i.e., a_t and b_t in $M_{t+1} = a_t + b_t R_{vw,t+1}$) will be a function of the consumption–aggregate wealth ratio. It follows that the consumption–aggregate wealth ratio may play a special role in both the CAPM and the CCAPM when they are specified as linear factor models with time-varying coefficients.

To show that the consumption–aggregate wealth ratio summarizes agents' expectations of future returns, make a loglinear approximation to a representative investor's intertemporal budget constraint, $W_{t+1} = (1 + R_{m,t+1})(W_t - C_t)$, and the log consumption–wealth ratio may be ex-

pressed in terms of future returns to the market portfolio and future consumption growth. (A full derivation of this approximation is given in App. A.) Because this approximate equation holds merely as a consequence of the agent's intertemporal budget constraint, it holds ex post, but it also holds ex ante. Accordingly, the log consumption–wealth ratio may be expressed as¹²

$$c_t - w_t \approx E_t \sum_{i=1}^{\infty} \rho_w^i (r_{m,t+i} - \Delta c_{t+i}). \quad (9)$$

Equation (9) implies that, if the conditional expectation of consumption growth is not too volatile (something that appears to be true empirically), the log consumption–wealth ratio summarizes expectations of future returns to the market portfolio. The specification in (9) is directly analogous to the linearized formula for the log dividend–price ratio (Campbell and Shiller 1988), where consumption enters in place of dividends and wealth enters in place of price. If the consumption–wealth ratio is high, then the agent must be expecting either high returns on wealth in the future or low consumption growth rates. The key difference between the consumption–wealth ratio and the dividend–price ratio is what is on the right-hand side: in (9) it is the return to the *entire* market portfolio and consumption growth; for the dividend–price ratio it is the return to the stock market component of wealth and dividend growth.

Of course, the log consumption–aggregate wealth ratio is not observable because human capital is not observable. To overcome this obstacle, Lettau and Ludvigson (2001) reformulate the bivariate cointegrating relation between c and w in (9) as a trivariate cointegrating relation involving three observable variables, namely log consumption, c , log nonhuman or asset wealth, a , and log labor earnings, y . Such a reformulation is possible under the condition that labor income is integrated and the rate of return to human capital is stationary. With these assumptions, the log of human capital, h_t , may be written as $h_t = \kappa + y_t + v_t$, where κ is a constant and v_t is a mean zero stationary random variable. This formulation for h_t may be rationalized by a number of different specifications linking labor income to the stock of human capital.¹³

With this assumption, we are now in a position to express the log consumption–aggregate wealth ratio in terms of observable variables.

¹² We omit unimportant linearization constants in linearized equations.

¹³ One such specification models aggregate labor income as the dividend paid from human capital, as in Campbell (1996) and Jagannathan and Wang (1996). In this case, the return to human capital is defined as $R_{h,t+1} = (H_{t+1} + Y_{t+1})/H_t$ and a loglinear approximation of $R_{h,t+1}$ implies that $v_t = E_t \sum_{j=0}^{\infty} \rho_h^j (\Delta y_{t+1+j} - r_{h,t+1+j})$.

Let A_t be nonhuman, or asset, wealth. Aggregate wealth is therefore $W_t = A_t + H_t$ and log aggregate wealth may be approximated as $w_t \approx \omega a_t + (1 - \omega)h_t$, where ω equals the average share of nonhuman wealth in total wealth, A/W . With this approximation, the left-hand side of (9) may be expressed as the difference between log consumption and a weighted average of log asset wealth and log labor income:

$$\begin{aligned} cay_t &\equiv c_t - \omega a_t - (1 - \omega)y_t \\ &\approx E_t \sum_{i=1}^{\infty} \rho_w^i (r_{m,t+i} - \Delta c_{t+i}) - (1 - \omega)v_t. \end{aligned} \quad (10)$$

Because all the variables on the right-hand side of (10) are stationary, the model implies that consumption, asset wealth, and labor income share a common stochastic trend (they are cointegrated), with ω and $1 - \omega$ parameters of this shared trend. As long as the last term on the right-hand side is not too variable, this equation implies that the observable quantity on the left-hand side should be a good proxy for the log consumption–aggregate wealth ratio and therefore expected returns. Lettau and Ludvigson (2001) present evidence that cay_t is a strong predictor of excess stock returns on aggregate stock market indexes.

An important task in using the left-hand side of (10) as a scaling variable is the estimation of the parameters in cay_t . Lettau and Ludvigson discuss how these parameters can be estimated consistently and why measurement considerations suggest that the coefficients on asset wealth and labor income may sum to a number less than one. The reader is referred to Lettau and Ludvigson (2001) for details on data construction and data definition and for a description of the procedure used to estimate ω and $1 - \omega$.¹⁴ We simply note here that we obtain an estimated value for cay_t , denoted $\widehat{cay}_t = c_t^* - 0.31a_t^* - 0.59y_t^* - 0.60$, where starred variables indicate measured quantities.¹⁵ We use this estimated value as a scaling variable in our empirical investigation.

¹⁴ The data used for the estimation of cay_t are quarterly, seasonally adjusted, per capita variables, measured in 1992 dollars. The consumption data pertain to nondurables and services excluding shoes and clothing in 1992 chain-weighted dollars. The nonhuman wealth data are the household net worth series provided by the Board of Governors of the Federal Reserve System. Labor income is defined as wages and salaries plus transfer payments plus other labor income minus personal contributions for social insurance minus taxes. Taxes are defined as [wages and salaries/(wages and salaries + proprietors income with IVA and Ccadj + rental income + personal dividends + personal interest income)] \times (personal tax and nontax payments), where IVA is inventory evaluation and Ccadj is capital consumption adjustments. Both the net worth variable and the labor income variable are deflated by the personal consumption expenditure chain-type price deflator. All variables are given in per capita terms.

¹⁵ The parameters in \widehat{cay}_t are estimated in a first-stage time-series analysis using only consumption, asset wealth, and labor income, and they are completely unrelated to the cross-sectional data on portfolio returns.

III. Econometric Specification and Tests

We use the beta representation (4) as the basis of our empirical work, specialized to the particular asset pricing model under consideration. In this section, we examine whether scaled specifications of the CAPM, the human capital CAPM, and the consumption CAPM (jointly, the (C)CAPM) can explain the cross section of expected returns with the conditioning variable, z_p , equal to \widehat{cay}_p . In each case, the scaled multi-factor beta representation (4) nests an associated unconditional model in which the β 's on the scaling variable and on the scaled factors are zero. We compare the ability of all these models to explain the cross section of average returns with that of the three-factor model in Fama and French (1993). In addition, we follow the suggestion in Jagannathan and Wang (1998) to include firm size (market equity) and book-to-market characteristics as additional explanatory variables as a test for model misspecification. We also test the (C)CAPM models using two alternative sets of portfolios that include scaled returns. We briefly discuss the results of using alternative scaling variables.

A. *Econometric Tests and Portfolio Data*

The unconditional model in (4) can be consistently estimated by the cross-sectional regression methodology proposed in Fama and MacBeth (1973), an approach we use here. In principle, other empirical procedures are available for testing the model in (4). We discuss the relative merits of alternative methods in Section IV and simply note here that the Fama-MacBeth procedure has important advantages for our application, in which we have only a moderate number of quarterly time-series observations (fewer than 150 in our data) but in which we require a reasonably large number of asset returns to test the model's cross-sectional implications.¹⁶

For summarizing the goodness of fit of each empirical specification, we report the average squared pricing errors across all 25 portfolios as well as the R^2 of the cross-sectional regression showing the fraction of the cross-sectional variation in average returns that is explained by each model.¹⁷ Although the cross-sectional R^2 is not a formal test of model specification, it is an informative summary statistic of how well each model fits the data, and it neatly illustrates the anomaly emphasized by Fama and French (1992) that the classic CAPM explains virtually none

¹⁶ The data on nonhuman wealth, a , and on subcomponents of the personal consumption expenditure that we use are available only on a quarterly basis.

¹⁷ This goodness of fit measure follows Jagannathan and Wang (1996) and is given by $[\text{Var}_c(\bar{R}_i) - \text{Var}_c(\bar{\epsilon}_i)] / \text{Var}_c(\bar{R}_i)$, where $\bar{\epsilon}_i$ is the average residual for portfolio i , Var_c denotes a cross-sectional variance, and variables with bars over them denote time-series averages.

of the cross-sectional variation in returns on these portfolios. We also test whether the coefficients λ in (4) are statistically different from zero. Following Shanken (1992), we report the standard errors of these coefficients corrected for sampling error that arises because the regressors $\hat{\beta}$ are estimated in a first-stage time-series regression for each $R_{i,t+1}$.¹⁸ Because Jagannathan and Wang (1998) show that the Fama-MacBeth procedure does not necessarily overstate the precision of the standard errors if conditional heteroskedasticity is present, we also report the conventional t -statistics.¹⁹

Our data on returns consist of 25 portfolios formed according to the same criteria as those used in Fama and French (1992, 1993). These data are value-weighted returns for the intersections of five size portfolios and five book-to-market equity (BE/ME) portfolios on the New York Stock Exchange, the American Stock Exchange, and NASDAQ stocks in Compustat. The portfolios are constructed at the end of June, and market equity is market capitalization at the end of June. The ratio BE/ME is book equity at the last fiscal year end of the prior calendar year divided by market equity at the end of December of the prior year. This procedure is repeated for every calendar year from July 1963 to June 1998. We refer the reader to the Fama-French articles cited above for details and data characteristics. We convert the returns to quarterly data producing a time series spanning the third quarter of 1963 to the third quarter of 1998, that is, 141 observations for each of the 25 portfolios.

As advocated by Ferson, Sarkissian, and Simin (2000), we demean the scaling variable in all the empirical investigations of this paper.

IV. Empirical Results

A. *The CAPM and the Fama-French Model*

1. Familiar Unconditional Models

Using returns on the 25 size and book-to-market sorted portfolios described above, we now examine the power of various beta representations to explain the cross section of average returns. Table 1 presents results of estimating the empirical specification in (4) for the CAPM,

¹⁸ The beta coefficients are estimated from a single multiple time-series regression for each asset. We use the entire sample to estimate the β 's; a rolling regression approach is not appropriate in quarterly data in which fewer than 150 time-series observations are available.

¹⁹ Note that standard errors do not need to be adjusted to account for the use of the generated regressor \widehat{ca}_y . This follows from the fact that estimates of the parameters in \widehat{ca}_y are "superconsistent," converging to the true parameter values at a rate proportional to the sample size T rather than proportional to \sqrt{T} as in ordinary applications (Stock 1987).

TABLE 1
FAMA-MACBETH REGRESSIONS USING 25 FAMA-FRENCH PORTFOLIOS: λ_j COEFFICIENT
ESTIMATES ON BETAS IN CROSS-SECTIONAL REGRESSION

ROW	CONSTANT	$\widehat{\alpha y}_i$	Factors _{<i>t+1</i>}				$\widehat{\alpha y}_i \cdot$ Factors _{<i>t+1</i>}		R^2 (\bar{R}^2)
			R_{vw}	Δy	SMB	HML	R_{vw}	Δy	
1	4.18 (4.47) (4.45)		-.32 (-.27) (-.27)					.01 -.03	
2	3.21 (3.37) (1.87)		-1.41 (-1.20) (-.67)	1.26 (3.42) (1.90)				.58 .54	
3	1.87 (1.31) (1.21)		1.33 (.83) (.76)		.47 (.94) (.86)	1.46 (3.24) (2.98)		.80 .77	
4	3.70 (3.88) (2.61)	-.52 (-.22) (-.15)	-.06 (-.05) (-.03)				1.14 (3.59) (2.41)	.31 .21	
5	3.70 (3.86) (2.60)		-.08 (-.07) (-.44)				1.16 (3.58) (2.41)	.31 .25	
6	5.18 (5.59) (3.32)	-.44 (-1.60) (-.95)	-1.99 (-1.73) (-1.02)	.56 (2.12) (1.26)			.34 (1.67) (.99)	-.17 (-2.40) (-1.42)	.77 .71
7	3.81 (4.02) (2.80)		-2.22 (-1.88) (-1.31)	.59 (2.20) (1.53)			.63 (2.79) (1.94)	-.08 (-2.52) (-1.75)	.75 .70

NOTE.—The table presents λ estimates from cross-sectional Fama-MacBeth regressions using returns of 25 Fama-French portfolios: $E[R_{i,t+1}] = E[R_{0,t}] + \beta \lambda$. The individual λ_j estimates (from the second-pass cross-sectional regression) for the beta of the factor listed in the column heading are reported. In the first stage, the time-series betas β are computed in one multiple regression of the portfolio returns on the factors. The term R_{vw} is the return of the value-weighted CRSP index, Δy_{t+1} is labor income growth, and SMB and HML are the Fama-French mimicking portfolios related to size and book-market equity ratios. The scaling variable is $\widehat{\alpha y}$. The table reports the Fama-MacBeth cross-sectional regression coefficient; in parentheses are two t -statistics for each coefficient estimate. The top statistic uses uncorrected Fama-MacBeth standard errors; the bottom statistic uses the Shanken (1992) correction. The term R^2 denotes the unadjusted cross-sectional R^2 statistic, and \bar{R}^2 adjusts for the degrees of freedom.

the human capital CAPM, and the Fama-French three-factor model. The table reports the estimated λ coefficients, Shanken-corrected and uncorrected t -statistics for these coefficients, and R^2 statistics for the cross-sectional regression. Table 2 presents test results for the joint significance of the scaled factors in each model. Section IVC discusses average pricing errors for each model.

To form a basis for comparison, we begin by presenting results from a series of familiar unconditional models. Of these, the most familiar is the static CAPM, with the CRSP value-weighted return, R_{vw} , used as a proxy for the unobservable market return. This model implies a cross-sectional specification taking the form

$$E[R_{i,t+1}] = E[R_{0,t}] + \beta_{vwi} \lambda_{vw} \quad (11)$$

The results are presented in the first row of table 1. The t -statistic for λ_{vw} shows that the beta on the value-weighted return is not a statistically significant determinant of the cross section of average returns. More-

TABLE 2
FAMA-MACBETH REGRESSIONS USING 25 FAMA-FRENCH PORTFOLIOS: TESTS FOR JOINT SIGNIFICANCE

Row	ALL	Factors _{<i>t</i>+1}	$\widehat{cay}_i \cdot \text{Factors}_{t+1}$	f_{t+1} AND $\widehat{cay}_i \cdot f_{t+1}$ FOR EACH FAC- TOR f	
				R_{vw}	Δy
1	.798				
2	.798				
	.000				
	.022				
3	.000				
	.002				
4	.000	.963	.000		
	.000	.975	.016		
5	.000	.948	.000		
	.003	.965	.016		
6	.000	.001	.000	.008	.000
	.000	.092	.021	.079	.040
7	.000	.001	.000	.001	.000
	.001	.032	.002	.012	.004

NOTE.—The table presents test results for the joint significance of the scaled factors in each model of table 1. Each row gives test statistics for the model in the corresponding row of table 1. p -values are reported from χ^2 tests of joint significance for four sets of variables: (i) all the right-hand-side betas, (ii) the betas for the fundamental factors, (iii) the betas for the scaled factors, and (iv) the beta for each fundamental factor joint with the beta for that factor scaled. The top number is computed using the uncorrected variance-covariance matrix, and the bottom number uses the Shanken (1992) correction. The model is estimated using data from 1963:Q3 to 1998:Q3. The coefficient estimates of the factors are multiplied by 100; the estimates of the scaled terms are multiplied by 1,000.

over, it has the wrong sign according to the CAPM theory. The R^2 for this regression summarizes this failure: only 1 percent of the cross-sectional variation in average returns can be explained by the beta for the market return. Note that the R^2 adjusted for degrees of freedom, denoted \bar{R}^2 , is negative. These results are now familiar (see Fama and French 1992). By contrast, a specification that includes—in addition to the value-weighted return beta—the beta for labor income growth advocated by Jagannathan and Wang (1996), Δy , performs much better, explaining about 58 percent of the cross-sectional variation in returns (row 2).²⁰

Row 3 of table 1 presents results for the Fama-French three-factor model given by

$$E[R_{i,t+1}] = E[R_{0,t}] + \beta_{vwi}\lambda_{vw} + \beta_{SMBi}\lambda_{SMB} + \beta_{HMLi}\lambda_{HML}, \quad (12)$$

where the “small minus big” (SMB) and “high minus low” (HML) port-

²⁰ Here we use the measure of labor income growth advocated by Jagannathan and Wang (1996): the growth in total personal, per capita income less dividend payments from the National Income and Product Accounts published by the Bureau of Economic Analysis. In addition, we followed the timing convention of Jagannathan and Wang (1996), in which labor income is lagged one month to capture lags in the official reports of aggregate income.

folios are constructed as in Fama and French (1993).²¹ In sharp contrast to the CAPM, this model explains about 80 percent of the cross-sectional variation in these returns, and the t -statistic on the HML factor is highly statistically significant even after correction for sampling error in the β 's. These results are consistent with what has been reported in the literature using monthly data. With these benchmark results in hand, we are now in a position to see how they change when they are modified to allow for conditioning information in \widehat{cay}_t .

2. Scaled Factor Models

Row 4 of table 1 shows results for the scaled, conditional CAPM with one fundamental factor, $f_{t+1} = R_{vw,t+1}$, and a single scaling variable $z_t = \widehat{cay}_t$. The cross-sectional regression for this scaled model takes the form

$$E[R_{i,t+1}] = E[R_{0,t}] + \beta_{zi}\lambda_z + \beta_{vwi}\lambda_{vw} + \beta_{vwi}\lambda_{vwz}. \quad (13)$$

The estimated value of λ_z is not statistically different from zero, implying that the time-varying component of the intercept is not an important determinant of average returns. Moreover, row 5 of table 1 shows that eliminating β_{zi} as an explanatory variable does not have an important effect on the marginal predictive power of the remaining betas or on the overall fit of the regression. We found this to be generally true in a variety of cross-sectional regressions we report in tables 3 and 4 below.

By contrast, the coefficient on β_{vwi} is strongly significant and is jointly significant with the coefficient on β_{vwi} (table 2).²² The R^2 statistic in row 4 for the scaled CAPM is considerably higher than for the simple static CAPM; it jumps to 31 percent from 1 percent by simply including β_{vwz} as an additional regressor.

The scaled CAPM specification including human capital takes the form

$$E[R_{i,t+1}] = E[R_{0,t}] + \beta_{zi}\lambda_z + \beta_{vwi}\lambda_{vw} + \beta_{vwi}\lambda_{vwz} + \beta_{\Delta yi}\lambda_{\Delta y} + \beta_{\Delta yzi}\lambda_{\Delta yz}. \quad (14)$$

The estimation results are presented in table 1, rows 6 and 7 (with and without β_{zi}). This model performs much better than the unscaled version and explains 75 percent of the cross-sectional variation in average returns, about as much as the Fama-French three-factor model in row 3.

²¹ The SMB portfolio is the difference between the returns on small and big stock portfolios with the same weighted-average book-to-market equity. The HML portfolio is the difference between returns on high- and low-BE/ME portfolios with the same weighted-average size. Further details on these variables can be found in Fama and French (1993).

²² These tests are carried out by forming a Wald statistic using either the uncorrected or the Shanken-corrected coefficient covariance matrix provided by the Fama-MacBeth procedure. Results for both covariance matrices are presented in table 2.

In particular, the coefficients on the scaled factors, λ_{vuz} and $\lambda_{\Delta yz}$, are statistically different from zero according to the uncorrected t -statistics. The Shanken correction to the t -statistics is large, however, especially so for models that include scaled macroeconomic factors rather than unscaled returns. The Shanken correction is negligible in the static CAPM, where the single factor is R_{vuz} , as exhibited in row 1 of table 1. The Shanken correction is directly related to the magnitude of each λ coefficient estimate and inversely related to factor variability. Thus, although the models with macro factors have smaller λ estimates than models with financial indicators as factors, the estimates of λ are not proportionally smaller relative to their smaller factor variance. We find these differences in the magnitude of the Shanken correction across models with and without macro variables to be a feature of all our tests, consistent with what has been found in other studies that include macroeconomic variables as factors (see, e.g., Shanken's [1992] example using macro data and Jagannathan and Wang [1996]).

Several other features of the cross-sectional results in table 1 bear noting. First, in the case of the CAPM (eq. [13]), the average risk price for the value-weighted return from the associated conditional linear factor model will be a weighted average of λ_{vw} and λ_{vuz} , the latter multiplied by z_i (see eq. [6]). Given these estimates and under the assumption that $\text{Var}_t(R_{vw,t+1})$ is approximately constant, the average risk price for the value-weighted return is found to be positive (recall the discussion in Sec. IIA). For the scaled human capital CAPM (eq. [14]), this same calculation yields a positive average risk price on the human capital beta. A problem with this model, however, is that there is a negative average risk price on the beta for the value-weighted return.²³

Second, the estimated value of the average zero-beta rate is large. We find this to be a feature of all the scaled models we test. The average zero-beta rate should be between the average "riskless" borrowing and lending rates, and the estimated value is implausibly high for the average investor. Although the (C)CAPM can explain a substantial fraction of the cross-sectional variation in these 25 portfolio returns, this result suggests that the scaled models do a poor job of simultaneously pricing the hypothetical zero-beta portfolio. This finding is not uncommon in studies that use macro variables as factors. For example, the estimated values for the zero-beta rate we find here have the same order of magnitude as that found in Jagannathan and Wang (1996). It is possible that the greater sampling error we find in the estimated betas of the scaled models with macro factors is contributing to an upward bias in the zero-beta estimates of those models relative to the estimates for

²³ Jagannathan and Wang (1996) report a similar finding for the signs of the risk prices on the market and human capital betas.

TABLE 3
 CONSUMPTION CAPM: FAMA-MACBETH REGRESSIONS USING 25 FAMA-FRENCH
 PORTFOLIOS
 A. λ_j COEFFICIENT ESTIMATES OF BETAS IN CROSS-SECTIONAL REGRESSIONS

Row	Constant	\widehat{cay}_t	Δc_{t+1}	$\widehat{cay}_t \cdot \Delta c_{t+1}$	R^2 (\bar{R}^2)
1	3.24		.22		.16
	(4.93)		(1.27)		.13
2	4.28	-.13	.02	.06	.70
	(6.10)	(-.43)	(.20)	(3.12)	.66
3	4.10		-.02	.07	.69
	(6.82)		(-.14)	(3.20)	.66
	(5.14)		(-.10)	(2.41)	

B. TESTS FOR JOINT SIGNIFICANCE				
Row	All	Δc_{t+1}	$\widehat{cay}_t \cdot \Delta c_{t+1}$	Δc_{t+1} and $\widehat{cay}_t \cdot \Delta c_{t+1}$
1	.205			
	.249			
2	.000	.840	.002	.000
	.001	.888	.030	.009
3	.000	.893	.001	.000
	.001	.919	.016	.001

NOTE.—See note to table 1. Δc denotes consumption growth.

models with only financial factors. Such arguments for large zero-beta estimates have a long tradition in the cross-sectional asset pricing literature (e.g., Black et al. 1972; Miller and Scholes 1972). However, if the sampling error is not independently and identically distributed and if multiple factors are poorly measured, there is little that can be said about the direction of bias. Procedures for discriminating the sampling error explanation for these large estimates of the zero-beta rate from others are not obvious, and its development is left to future research.

B. The Consumption CAPM

We now turn to a cross-sectional empirical analysis of the consumption CAPM. Table 3 presents, for the CCAPM, the same results presented in tables 1 and 2 for the CAPM and human capital CAPM. The scaled multifactor consumption CAPM, with $z_t = \widehat{cay}_t$ as the single conditioning variable, is a special case of (4) and takes the form

$$E[R_{i,t+1}] = E[R_{0,t}] + \beta_{zi}\lambda_z + \beta_{\Delta ci}\lambda_{\Delta c} + \beta_{\Delta czi}\lambda_{\Delta cz}, \quad (15)$$

where Δc denotes the log difference in consumption, as measured in

Lettau and Ludvigson (2001).²⁴ Thus the factors in this model are lagged \widehat{cay} , current-period consumption growth, and consumption growth times lagged \widehat{cay} . For comparison, row 1 of table 3 reports the unconditional consumption CAPM estimates (with consumption growth the single factor). The unconditional CCAPM performs only slightly better than the static CAPM, explaining just 16 percent of the cross-sectional variation in average returns on these 25 portfolios. The results of estimating the scaled specification are presented in row 2.

Row 2 shows that the time-varying component of the intercept term in (15) is not important: λ_z is not statistically significantly different from zero, and eliminating β_{zi} from the cross-sectional regression does not have an important effect on the other regression coefficients or on the overall regression fit (the table 1 results for the scaled CAPM models are comparable). This is illustrated in row 3, which presents results for the case in which β_{zi} is eliminated as a regressor in (15). In the interest of parsimony, from here on we focus on specifications in which the time-varying component of the constant is omitted (i.e., λ_z is restricted to be zero). None of the results presented in table 1 or table 4 below are qualitatively influenced by imposing this restriction.²⁵

Row 2 shows that the estimated values of $\lambda_{\Delta c}$ and $\lambda_{\Delta cz}$ are strongly jointly significant (panel B). (Because the scaled factor is included to allow for the possibility of time variation in the coefficient on that fundamental factor, there is no implication that the betas for the scaled and unscaled fundamental factor must be individually significant, only that they be jointly significant.) In addition, the coefficient on the scaled consumption factor $\lambda_{\Delta cz}$ is strongly individually significant (the uncorrected t -statistic for $\lambda_{\Delta cz}$ is 3.2; the corrected t -statistic is 2.41). More strikingly, the R^2 statistic indicates that the specification in (15) explains 70 percent of the cross-sectional variation in average returns on the 25 Fama-French portfolios. This result stands in sharp contrast to the 1 percent explained by the static CAPM (row 1, table 1) or the 16 percent explained by the unconditional consumption CAPM (row 1, table 3).

²⁴ Breeden et al. (1989) emphasize that measured quarterly consumption may be the time average of instantaneous consumption rates during the quarter. They show that one can compensate for this bias by multiplying quarterly consumption growth rates by $\frac{3}{4}$. Such an adjustment would scale the point estimates of the risk prices we estimate by $\frac{3}{4}$ but would obviously not effect the t -statistics, R^2 statistics, or average pricing errors we report.

²⁵ By contrast, results are somewhat sensitive to whether \widehat{cay} is included in the first-stage time-series regression. Note that the parameter b_j for a factor j may be nonzero in a pricing kernel with N factors, $M_{t+1} = a + \sum_{j=1}^N b_j f_{j,t+1}$, even if its beta is not priced (i.e., $\lambda_j = 0$ in the cross-sectional regression). This seems to be the case in our application: including the scaling variable \widehat{cay} as a factor in the pricing kernel is important even when the beta for this factor was not typically priced. Thus we always include the scaling variable as a factor in our specification of M_{t+1} .

Furthermore, it is quite close to the 80 percent R^2 produced from the Fama-French three-factor model (row 3, table 1).²⁶

Our finding that the beta for the cross term, $z_t \Delta c_{t+1}$, is important raises the possibility that the data may be well described by the habit models discussed above, for which this term is clearly an important part of the pricing kernel. Notice also that the improvement in the R^2 statistic obtained by scaling the unconditional CCAPM is quite dramatic and is larger than that derived from scaling the version of the CAPM that includes labor income growth. Finally, we note that, even though the coefficient $\lambda_{\Delta c}$ in (15) is negative, the implied average risk price for consumption risk in the underlying conditional consumption CAPM, $E[\tilde{\lambda}_{\Delta c,t}]$, computed as described above, is positive, consistent with the theory.

C. Average Pricing Errors

Figure 1, discussed in the Introduction, provides a visual impression of the relative empirical performance of each model we investigate. In this subsection we take a closer look at this figure. For a given empirical specification, we plot the fitted expected return for each of the 25 portfolios against their realized average returns. For reference, the data for these plots (the pricing errors for each portfolio in each empirical specification) are given in table 4. Each two-digit number in figure 1 represents one portfolio. The first digit refers to the size quintiles (1 indicating the smallest firms, 5 the largest), and the second digit refers to book-to-market quintiles (1 indicating the portfolio with the lowest book-to-market ratio, 5 the highest). Figure 1*a* plots fitted returns for the static CAPM, figure 1*b* plots fitted returns for the Fama-French model, and figures 1*c* and 1*d* do so for the unconditional CCAPM and the scaled CCAPM, respectively.

As mentioned, figure 1*a* shows that the simple, static CAPM explains virtually none of the variation in average returns on these portfolios. The main source of difficulty is immediately apparent: the mispricing of portfolios that have different book-to-market equity ratios for a given size value. For example, portfolios 11 (small growth) and 15 (small value)—those that are in the smallest size category but in the lowest

²⁶ It is not surprising that the Fama-French three-factor model explains a slightly larger fraction of the variation in average returns on these portfolios. If there is any measurement error in a set of theoretically derived aggregate indicators determining the discount factor, M , the factor mimicking portfolios for those variables will always price assets better than the underlying economic indicators. Mimicking portfolios are typically better measured and are often available on a more timely basis than macroeconomic data. On the other hand, if the Fama-French factors are not mimicking portfolios for consumption risk but are merely ex post mean-variance efficient portfolios, they will again always beat the theoretically derived factors in the sample.

TABLE 4
PRICING ERRORS

Portfolio	HC-CAPM					
	CAPM	CAPM Scaled	CCAPM	CCAPM Scaled	Scaled	Fama-French
A. Individual Portfolios						
S1B1	-1.0378	-1.1857	-1.7303	-1.0083	-.6451	-.8611
S1B2	.2601	.0859	-.3686	-.1109	-.0435	.0055
S1B3	.3494	.1532	-.0105	-.0809	-.1852	-.0276
S1B4	.9048	1.0021	.4774	.3803	.5386	.3722
S1B5	1.3108	.7976	.8278	.3704	.3504	.1555
S2B1	-.7860	-.7558	-1.0749	-.0362	-.1149	-.1416
S2B2	-.0189	.3796	-.0847	-.2823	-.2661	.0478
S2B3	.6786	.8639	.5696	.5310	.7516	.5134
S2B4	.7561	.6316	.7800	.0964	.0753	.3244
S2B5	1.0436	.4944	.7576	.3070	-.0309	.0757
S3B1	-.8062	-.2985	-.7324	.3645	-.0300	-.0310
S3B2	-.0659	.1049	.1019	-.0868	.1712	.2053
S3B3	.0486	.2463	.1996	-.4457	-.0357	-.0291
S3B4	.4098	.0956	.6398	.2218	-.1837	.1066
S3B5	.9036	1.1403	.8631	.4109	.7183	.0225
S4B1	-.6513	-.1532	-.3754	.4252	.1218	.4184
S4B2	-.8487	-.2545	-.6870	-.3924	-.4210	-.5041
S4B3	-.1903	-.2714	.1383	-.0626	-.2180	-.1749
S4B4	.2339	-.3943	.4651	.0130	-.3444	-.0347
S4B5	.7553	-.2413	.6952	.3610	.0678	-.0073
S5B1	-.7959	-.9014	-.4572	.4159	.6455	.5342
S5B2	-.8186	-.7850	-.2918	-.7343	-.2759	.0035
S5B3	-.9287	-.3369	-.6240	-.4895	-.1888	-.1980
S5B4	-.4252	-.2523	.0760	.1226	-.2116	-.2715
S5B5	-.2811	-.1650	-.1547	-.2901	-.2457	-.5043
B. Pricing Errors of Aggregated Portfolios						
S1	.872	.784	.899	.513	.416	.421
S2	.740	.649	.731	.305	.362	.257
S3	.573	.542	.590	.334	.341	.094
S4	.601	.274	.516	.306	.269	.333
S5	.697	.572	.378	.458	.356	.336
B1	.825	.762	1.004	.549	.415	.484
B2	.541	.411	.377	.398	.267	.253
B3	.545	.451	.393	.382	.370	.239
B4	.599	.571	.542	.209	.314	.276
B5	.924	.673	.709	.351	.375	.211
Average	.705	.589	.648	.393	.352	.308
χ^2	63.67*	27.24	53.03*	33.88	27.54	45.33*

NOTE.—This tables reports the pricing errors (in percent) from the Fama-MacBeth regressions presented in tables 1–3. Panel A lists the average errors of each Fama-French portfolio for the model listed in the column heading; HC-CAPM refers to the human capital CAPM described in the text, S1 refers to the portfolios with the smallest firms, and S5 includes the largest firms. Similarly, B1 includes firms with the lowest book-to-market ratio and B5 the highest. Panel A reports pricing errors for the 25 size and book-to-market sorted portfolios, and panel B computes the square root of the average squared pricing errors for aggregated portfolios. The last two rows report the square root of the average squared pricing errors across all portfolios and a χ^2 statistic for a test that the pricing error is zero. The model is estimated using data from 1963:Q3 to 1998:Q3.

* Statistically different from zero at the 5 percent level.

and highest book-to-market categories—lie farthest from the 45-degree line. Similar mispricing occurs for portfolios 21 and 25 (semismall growth and semismall value), 31 and 35 (medium-sized growth and medium-sized value), and 41 and 45 (semilarge growth and semilarge value). For each of these pairs of portfolios, the value portfolio earns a higher average return than the growth portfolio, yet the CAPM predicts that they should earn roughly the same expected return. Thus the figure illustrates a well-known result: it is this “value effect” that destroys the static CAPM when confronted with these portfolios. The contrast with the Fama-French model in figure 1*b* is stark.

Figure 1*c* shows that the unscaled consumption CAPM also has difficulty explaining the difference in return between high- and low-book-to-market portfolios. Again the returns on portfolios 15 and 25 are substantially higher than those on portfolios 11 and 21, respectively, yet the fitted expected returns for each of these two pairs of portfolios are roughly the same. By contrast, the scaled multifactor consumption CAPM does a much better job of explaining the value effect (fig. 1*d*): the fitted expected returns on value portfolios are high whereas the fitted expected returns on growth portfolios are low, consistent with the data.

When the fitted returns from the scaled CCAPM are compared with those for the Fama-French model in figure 1*b*, it is evident that the scaled consumption CAPM does about as well as the Fama-French model in explaining this value effect. Also, the *pattern* of mispricing in the Fama-French model is very similar to that in the scaled CCAPM. For example, both models do a much better job of pricing portfolios 11 and 15, 21 and 25, 31 and 35, and 41 and 45 than the static CAPM. Both models have relatively greater difficulty pricing portfolios 11, 51, and 15 than they have pricing other portfolios.

How do the pricing errors vary across more aggregated portfolios? Panel B of table 4 reports the square root of the average squared pricing errors across 10 aggregated portfolios formed on the basis of the size and book-to-market quintiles. It is clear that the pricing errors for the scaled consumption CAPM are lower for large-size portfolios and high-book-to-market portfolios, and vice versa.²⁷ A similar pattern is displayed for the Fama-French model. The last row of table 4 gives the square root of the average squared pricing errors across all portfolios. The average squared pricing error for the conditional consumption CAPM is a little over half as large as that of the simple static CAPM; the average squared pricing error for the Fama-French model is about 42 percent

²⁷ These results are consistent with those of Avramov (1999), who finds that *cay* has important predictive power for returns on large and medium as well as on high-book-to-market portfolios in a Bayesian study of return forecasting models.

as large as that of the static CAPM. Notice that the average squared pricing errors display the same relative pattern across models that the R^2 statistics do. Thus, if models were ranked by their average pricing errors, the same ranking would be obtained using cross-sectional R^2 statistics.

Just below the average squared pricing errors are the results of an asymptotic χ^2 (Wald) test of the null hypothesis that all the pricing errors are jointly zero.²⁸ The table shows that the only models for which the null of zero pricing errors may not be rejected are the scaled multifactor models; the unscaled models, including the Fama-French model, are all statistically rejected according to this test. We are reluctant to place emphasis on this result. It is clear that the Fama-French three-factor model, the scaled human capital CAPM, and the scaled consumption CAPM all have average pricing errors of roughly the same magnitude, and the economic size of these errors is not large. The difference in statistical significance appears to be attributable to the greater sampling error in the estimated betas of the scaled macro factors. Such sampling error translates into a larger upward correction to the asymptotic variance-covariance matrix of the pricing errors. But, more significantly, several investigations have found that these tests, which rely on a consistent estimate of the variance-covariance matrix of pricing errors, have poor small-sample properties (e.g., Burnside and Eichenbaum 1996; Hansen, Heaton, and Yaron 1996). The small sample size of the Wald tests exceeds the asymptotic size, and this is especially true when the number of moment restrictions is large relative to the time-series sample size as it is in our empirical application.

D. Intuition

Why does the scaled CCAPM do so much better than the unscaled models at explaining why value stocks earn higher returns than growth stocks? This subsection provides some intuition for these findings. Per-

²⁸ The test statistic requires the assumption that the errors in the Fama-MacBeth regressions are independently and identically distributed over time. It is given by

$$(1 + \lambda' \Sigma_f^{-1} \lambda)^{-1} \hat{\alpha}'_{FM} \text{Cov}(\hat{\alpha}_{FM})^{-1} \hat{\alpha}_{FM} \sim \chi^2_{N-K}$$

where Σ_f is the variance-covariance matrix of the factors, $\hat{\alpha}_{FM}$ is the estimated vector of pricing errors given by the Fama-MacBeth estimates, N is the number of portfolios, and K is the number of factors. The first multiplicative term in parentheses is a correction for sampling error in β and is due to Shanken (1992). Note that this formula is numerically equivalent to the analogous test statistic for ordinary least squares (OLS) estimates from cross-sectional regressions, where the OLS standard errors are corrected for cross-sectional correlation:

$$T(1 + \lambda' \Sigma_f^{-1} \lambda)^{-1} \hat{\alpha}'_{OLS} [(I_N - \beta(\beta'\beta)^{-1}\beta') \Sigma [I_N - \beta(\beta'\beta)^{-1}\beta']]^{-1} \hat{\alpha}_{OLS} \sim \chi^2_{N-K}$$

where Σ is the $N \times N$ covariance matrix of pricing errors from the OLS cross-sectional regression (see Cochrane 2001, chap. 12).

haps this intuition can be most easily grasped by referring back to equation (8). This equation emphasizes that an asset's risk is determined, not by a simple correlation of its return with consumption growth, but by that correlation *conditional* on some state variable that reflects time variation in risk premia. This time variation in risk premia may be attributable to time variation in risk aversion (as in models with habit persistence, e.g., Campbell and Cochrane [1999]) or time variation in risk itself (as in models with time-varying labor earnings or default risk, e.g., Constantinides and Duffie [1996] and Chang and Sundaresan [1999]). Either way, the concept of risk is very different than in unconditional models such as the static CAPM or the unconditional consumption CAPM; it pertains to conditional rather than unconditional correlations with consumption growth.

If this conditionality is important empirically, it should be well captured by scaling consumption growth with \widehat{cay} , a proxy for the consumption-wealth ratio. Lettau and Ludvigson (2001) find that \widehat{cay} is a powerful forecaster of excess returns on aggregate stock market indexes; therefore, when \widehat{cay} is high, risk premia are expected to rise, and when \widehat{cay} is low, they are expected to fall.

Accordingly, the results above suggest that value portfolios are riskier than growth stocks not because their returns are more highly unconditionally correlated with consumption growth, but because their returns are more highly correlated with consumption growth when risk/risk aversion is high (\widehat{cay} is high) than when risk/risk aversion is low (\widehat{cay} is low). We explore this possibility in figure 2. Figure 2 plots, for pairs of value and growth portfolios in the same size category, the consumption beta conditional on being in a "bad" state (a period of high risk/risk aversion or high \widehat{cay}) along with the consumption beta conditional on being in a "good" state (a period of low risk/risk aversion or low \widehat{cay}). These conditional consumption betas are obtained from the following time-series regression, which is implied by our model for each portfolio return, R_t^i :

$$R_{t+1}^i = \alpha + \beta_{\Delta c}^i \Delta c_{t+1} + \beta_{\Delta cz}^i \Delta c_{t+1} z_t + \beta_{z_p}^i z_p$$

where in our empirical specification, the scaling variable $z_t = \widehat{cay}_t$. Collecting terms on Δc_{t+1} , we define a conditional consumption beta for the i th portfolio as $B_t^i \equiv \beta_{\Delta c}^i + \beta_{\Delta cz}^i z_t$. This says that the conditional correlation of each portfolio's return with consumption growth is a function of \widehat{cay} . Figure 2 plots the average conditional consumption beta for portfolio i in state s , where $s = \text{bad, good, equal}$ to $B_s^i \equiv \beta_{\Delta c}^i + \beta_{\Delta cz}^i \bar{z}_s$, where \bar{z}_s is the average value of \widehat{cay} in state s . For this exercise, a good state is defined as a quarter during which \widehat{cay}_t is at least one standard deviation below its mean, and a bad state is a quarter during which \widehat{cay}_t is at least one standard deviation above its mean. To avoid clutter,

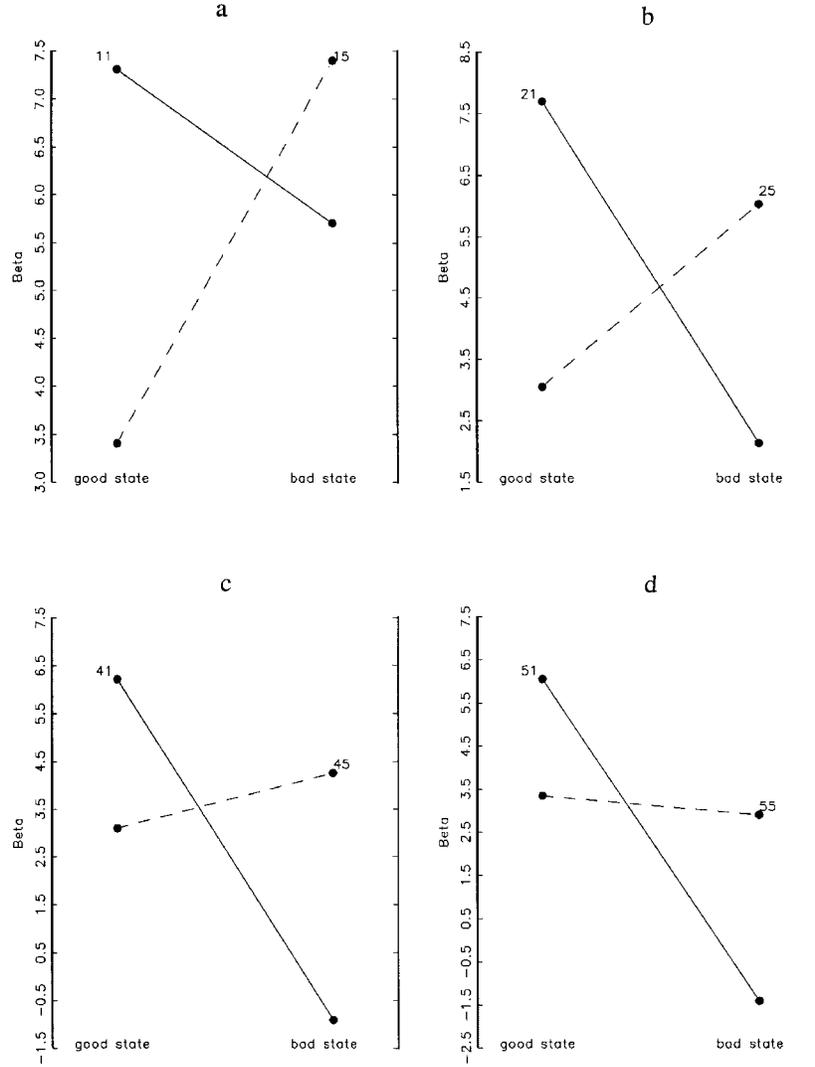


FIG. 2.—Conditional consumption betas in good and bad states. Each part of the figure displays the average consumption beta in good and bad states, conditional on z_{t-1} for portfolio i , equal to $B_s^i \equiv \beta_{\Delta c}^i + \beta_{\Delta z}^i \bar{z}_s$, where \bar{z}_s is the average value of \widehat{cay} in state s , $s =$ good or bad. A good state is defined as a quarter in which the scaling variable \widehat{cay} is one standard deviation below its mean value. A bad state is defined as a quarter in which the scaling variable \widehat{cay} is one standard deviation above its mean value. Part *a* compares B_s^i for the small-growth portfolio 11 and the small-value portfolio 15 (portfolios in the smallest size category and smallest and largest book-market category, respectively); part *b* compares B_s^i for the semismall-growth portfolio 21 and the semismall-value portfolio 25 (portfolios in the next-to-smallest size category and smallest and largest book-market category, respectively); part *c* compares B_s^i for the semilarge-growth portfolio 41 and the semilarge-value portfolio 45 (in the next-to-largest size category and smallest and largest book-market category, respectively); part *d* compares B_s^i for the large-growth portfolio 51 and the large-value portfolio 55 (in the largest size category and smallest and largest book-market category, respectively).

TABLE 5
CONDITIONAL BETAS IN CONSUMPTION CAPM

Portfolio	All States	Good States	Bad States
S1B1	6.3529	7.3089	5.7052
S1B2	6.3032	5.8321	6.6223
S1B3	5.1117	4.4374	5.5686
S1B4	5.5011	4.4285	6.2277
S1B5	5.7885	3.4071	7.4018
S2B1	4.3819	7.6970	2.1361
S2B2	3.6303	3.6265	3.6329
S2B3	3.9157	4.1440	3.7611
S2B4	3.5227	2.0702	4.5068
S2B5	4.8271	3.0495	6.0314
S3B1	2.7042	7.3587	-.4491
S3B2	2.7613	3.4644	2.2850
S3B3	2.9171	1.9948	3.5419
S3B4	2.5753	2.5729	2.5770
S3B5	3.7081	2.7064	4.3867
S4B1	1.9670	6.2181	-.9130
S4B2	2.6187	4.7651	1.1647
S4B3	1.9439	2.9882	1.2365
S4B4	2.5048	2.3777	2.5908
S4B5	3.7845	3.0944	4.2520
S5B1	1.6045	6.0536	-1.4096
S5B2	1.1598	2.0407	.5629
S5B3	2.3243	4.0597	1.1487
S5B4	1.2386	3.4376	-.2511
S5B5	3.0729	3.3378	2.8935

NOTE.—The table reports the average consumption betas in good and bad states, conditional on z_{t-1} for portfolio i , equal to $B_i^s = \beta_{i,c}^s + \beta_{i,c}^s \bar{z}_i$, where \bar{z}_i is the average value of \widehat{cay}_i in state s , $s = \text{good or bad}$. A good state is defined as a quarter in which the scaling variable \widehat{cay}_i is one standard deviation below its mean value. A bad state is defined as a quarter in which the scaling variable \widehat{cay}_i is one standard deviation above its mean value.

figure 2 plots only those conditional betas for the portfolios that are most mispriced by the static CAPM. The estimated conditional betas for the full set of 25 portfolios are given in table 5.

If time variation in risk premia is important for explaining why value portfolios are risky, those portfolios should have higher consumption betas in bad times (when \widehat{cay}_i is high) than in good times (when \widehat{cay}_i is low). In addition, if value portfolios are riskier than growth portfolios, they should have higher consumption betas in bad times than growth stocks. Figures 2a–2c show that this is precisely what is found. Value portfolios (defined to be a portfolio that is in the highest book-market category for a given size category) have higher consumption betas in bad states than the growth portfolios in the same size category (a growth portfolio is defined as one in the lowest book-market category for a given size category). Thus $B_{\text{bad}}^{15} > B_{\text{bad}}^{11}$, $B_{\text{bad}}^{25} > B_{\text{bad}}^{21}$, and $B_{\text{bad}}^{45} > B_{\text{bad}}^{41}$. By contrast, growth portfolios have higher consumption betas in good states than value portfolios, that is, $B_{\text{good}}^{15} < B_{\text{good}}^{11}$, $B_{\text{good}}^{25} < B_{\text{good}}^{21}$, and $B_{\text{good}}^{45} < B_{\text{good}}^{41}$. Furthermore, value portfolios are more highly correlated with

consumption in bad times than they are in good times, that is, $B_{\text{bad}}^{15} > B_{\text{good}}^{15}$, $B_{\text{bad}}^{25} > B_{\text{good}}^{25}$, and $B_{\text{bad}}^{45} > B_{\text{good}}^{45}$; the opposite is true for growth portfolios, that is, $B_{\text{bad}}^{11} < B_{\text{good}}^{11}$, $B_{\text{bad}}^{21} < B_{\text{good}}^{25}$, and $B_{\text{bad}}^{41} < B_{\text{good}}^{45}$. The semilarge growth portfolio 41 even provides a bit of consumption insurance, yielding a slight negative covariance with consumption growth in bad states. Thus value stocks are more highly correlated with consumption growth during times in which the representative investor least wants them to be, thereby making them riskier than growth stocks. Unconditional CCAPM models fail to capture the state dependency displayed by the criss-cross pattern of consumption betas in figure 2 because they assume that each asset's covariance with consumption growth is constant. This explains why the conditional consumption CAPM does a much better job than unconditional models of pricing these portfolios, correctly predicting that portfolios 15, 25, and 45 have higher average returns than portfolios 11, 21, and 41, respectively. The static CAPM and unconditional consumption CAPM miss this conditionality in the correlation of returns with consumption growth and, as a consequence, tend to price these portfolios poorly.

Figure 2*d* illustrates this intuition from a slightly different perspective. An inspection of figures 1*a* and 1*c* shows that the large-value portfolio 55—unlike the other portfolios plotted in figures 2*a*–2*c*—is almost perfectly priced by both the static CAPM and the unconditional consumption CAPM. The reason for this is that the consumption beta for portfolio 55 is about the same in good states as it is in the bad states ($B_{\text{good}}^{55} \approx B_{\text{bad}}^{55}$). Because this portfolio's correlation with consumption growth just happens to be independent of \widehat{cay}_b , the static CAPM and unconditional consumption CAPM price its average return well.

Table 5 shows that the general pattern displayed in figure 2—namely, that value stocks are more highly correlated with consumption growth in bad times than they are in good times and that value stocks are more highly correlated with consumption growth in bad times than growth stocks—holds for all five pairs of value-growth portfolios in a given size category.

Another way to provide intuition for our findings is to follow the novel approach taken in Lakonishok, Shleifer, and Vishny (1994). They reason that value stocks would be fundamentally riskier than growth stocks if (i) the former underperforms the latter in some states of the world and (ii) these times of underperformance are, on average, bad states rather than good states. Figure 3 shows the quarterly excess return on a value over a growth stock in the same size category. The excess returns displayed pertain to four pairs of portfolios that are very poorly priced by the static CAPM but are well priced by the Fama-French model and scaled CCAPM. These are the excess returns on portfolio 15 over 11, 25 over 21, 35 over 31, and 45 over 41.

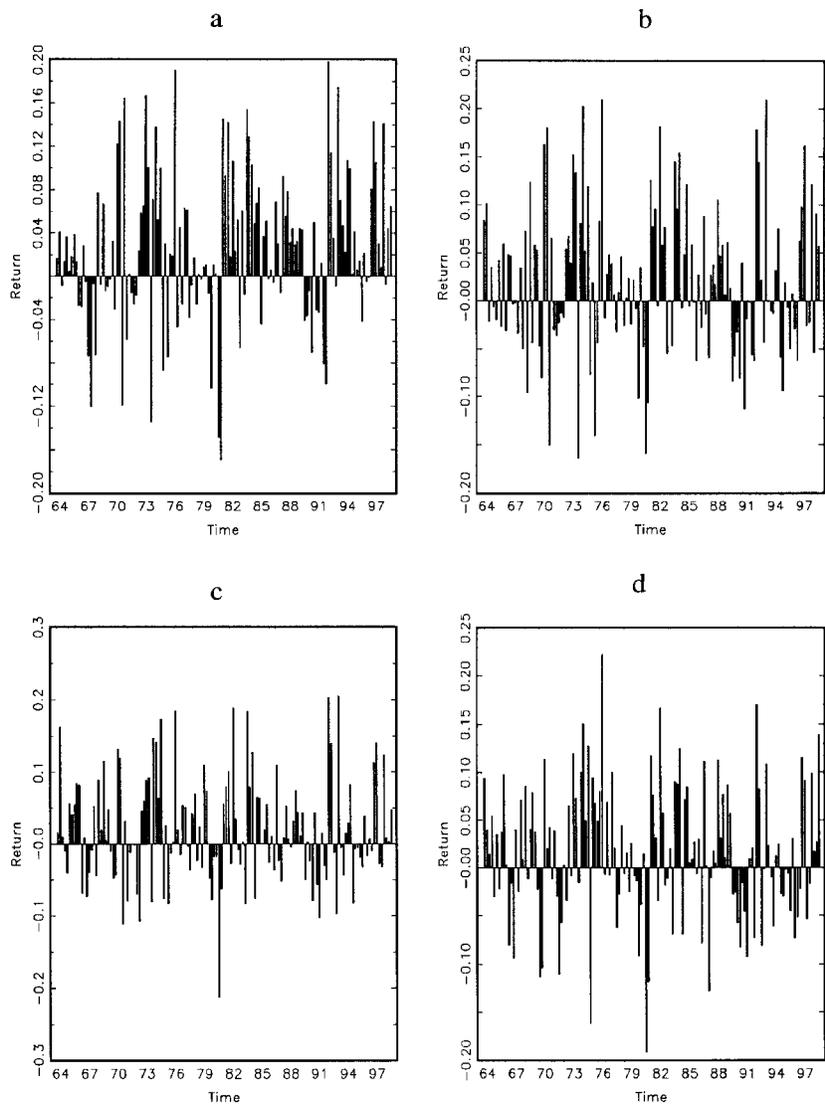


FIG. 3.—Value minus growth returns. This figure shows the quarterly excess return on a value over a growth stock in the same size category: The excess returns displayed are for portfolio 15 over 11 in part *a*, 25 over 21 in part *b*, 35 over 31 in part *c*, and 45 over 41 in part *d*.

The four parts of figure 3 show that there are many quarters during which the value portfolio underperformed the growth portfolio in the same size category. In fact, underperformance in these four parts occurs between 35 and 50 percent of the time. Thus condition i is satisfied in our data. A simple calculation shows that, for each of these excess returns, periods of underperformance were quarters in which \widehat{cay} was, on average, above its mean (i.e., risk premia were above average), whereas quarters of overperformance were times in which \widehat{cay} was, on average, below its mean (i.e., risk premia were below average). Thus periods of underperformance tend to coincide with bad states, and condition ii is satisfied.

Although these conclusions are quite different from those reached by Lakonishok et al., there are several ways in which our analyses differ. Our definition of value and growth portfolios, the horizon over which we measure excess returns, and the variable we use to determine good and bad states all differ from those of Lakonishok et al. but are the appropriate choices for our empirical application. On the other hand, our results are not that surprising given the finding in Lakonishok et al. that the buy-and-hold horizon matters. The shortest horizon considered by Lakonishok et al. is one year, and as they consider longer horizons they find that value stocks more consistently outperform their definition of growth stocks. We find that value stocks underperform growth stocks frequently at quarterly buy-and-hold horizons, suggesting that important movements in conditional moments occur at horizons of less than one year.

E. Including Characteristics

This subsection investigates whether there are any residual effects of firm characteristics in the scaled (C)CAPM models investigated above. We do this because several authors have argued that such a procedure provides a test of model misspecification.²⁹ This examination is done by alternately including portfolio size—the time-series average of the log of market equity for each portfolio—and the portfolio book-market equity ratio as additional explanatory variables in the cross-sectional regressions. A large *t*-value on the characteristic term suggests that the model may be misspecified. These results are presented in table 6.

²⁹ Kan and Zhang (1999) argue that “useless” factors can appear statistically important in the Fama-MacBeth methodology when the model being tested is misspecified. However, Berk (1995) and Jagannathan and Wang (1998) show that this misspecification can be tested for by including firm-specific characteristics as additional explanatory variables in cross-sectional asset pricing tests. Jagannathan and Wang (1998, theorem 6) prove that a useless factor cannot drive out a firm characteristic in the cross-sectional (second-pass) regression.

TABLE 6
 FAMA-MACBETH REGRESSIONS INCLUDING CHARACTERISTICS
 A. λ_j ESTIMATES ON BETAS IN CROSS-SECTIONAL REGRESSIONS INCLUDING SIZE

ROW	CONSTANT	Factors _{t+1}			$\widehat{cay}_t \cdot$ Factors _{t+1}			SIZE	R^2 (\bar{R}^2)
		R_{tw}	Δy	Δc	R_{tw}	Δy	Δc		
1	14.18	-3.60						-.57	.70
	(4.77)	(-2.78)						(-3.46)	.67
	(4.35)	(-2.54)						(-3.15)	
2	13.10	-3.05			.82			-.49	.75
	(4.71)	(-2.49)			(3.14)			(-3.24)	.73
	(3.79)	(-2.01)			(2.52)			(-2.61)	
3	12.03	-3.00	.51					-.41	.74
	(4.56)	(-2.52)	(2.00)					(-2.81)	.70
	(3.73)	(-2.06)	(1.63)					(-2.30)	
4	10.33	-2.68	.33		.59	-.02		-.33	.80
	(3.78)	(-2.33)	(1.36)		(2.63)	(-.59)		(-1.93)	.76
	(2.97)	(-1.84)	(1.07)		(2.07)	(-.46)		(-1.52)	
5	5.59			.04				-.18	.22
	(2.04)			(.35)				(-1.11)	.15
	(2.03)			(.35)				(-1.10)	
6	6.09			-.16			.08	-.15	.72
	(2.21)			(-1.45)			(3.23)	(-.87)	.68
	(1.66)			(-1.09)			(2.42)	(-.65)	

B. λ_j ESTIMATES ON BETAS IN CROSS-SECTIONAL REGRESSIONS INCLUDING BOOK-MARKET RATIO

ROW	CONSTANT	Factors _{t+1}			$\widehat{cay}_t \cdot$ Factors _{t+1}			BOOK-MARKET RATIO	R^2 (\bar{R}^2)
		R_{tw}	Δy	Δc	R_{tw}	Δy	Δc		
1	2.25	1.47						1.17	.82
	(2.06)	(1.08)						(3.62)	.81
	(2.01)	(1.06)						(3.57)	
2	2.22	1.45			.15			1.12	.83
	(2.01)	(1.05)			(.77)			(3.51)	.81
	(1.95)	(1.02)			(.75)			(3.41)	
3	1.91	2.00	.41					1.38	.83
	(1.68)	(1.41)	(1.61)					(3.89)	.80
	(1.52)	(1.29)	(1.44)					(3.53)	
4	2.81	.97	-.23		.14	-.05		1.09	.85
	(2.56)	(.71)	(-.94)		(.70)	(-1.56)		(3.13)	.81
	(2.36)	(.66)	(-.86)		(.64)	(-1.44)		(2.88)	
5	3.69			.14				.83	.75
	(5.98)			(.81)				(2.81)	.73
	(5.70)			(.77)				(2.67)	
6	3.90			.08			.02	.61	.78
	(6.29)			(.55)			(1.40)	(1.86)	.75
	(5.95)			(.52)			(1.32)	(1.75)	

NOTE.—See notes to tables 1–3. This table presents estimates of cross-sectional Fama-MacBeth regressions using the returns on 25 Fama-French portfolios:

$$E[R_{i,t+1}] = E[R_{i,t}] + \beta\lambda + d\Theta,$$

where Θ denotes a characteristic variable: Θ is either the log of the portfolio size (size in panel A) or the log of the portfolio book-to-market ratio (in panel B).

The results in row 1 are comparable with what has been reported elsewhere for the static CAPM: the coefficient on size is strongly significant, and the R^2 statistic jumps from 1 to 70 percent when it is included as an explanatory variable. Also, the risk price for the value-weighted return is now negative and statistically significant. Once labor income growth is included in the scaled CAPM (row 4), these size effects are attenuated but not entirely eliminated; the coefficient on size is not statistically different from zero according to the corrected t -statistic but remains statistically significant according to the uncorrected t -statistic. By contrast, size is not a significant determinant of the cross section of average returns in either the unconditional consumption CAPM or the scaled consumption CAPM; the coefficient on this variable is not statistically significant, and the overall fit of the regression is roughly the same regardless of whether size is included in the regression.

Panel B of table 6 shows that the book-market ratio is also highly significant when included in both the unscaled and scaled versions of the CAPM, and the R^2 statistic increases by more than 80 percent once the book-market ratio is included in the static CAPM model. Even the human capital CAPM has difficulty eliminating residual book-to-market effects regardless of whether this model is scaled (rows 3 and 4). A similar result holds for the unscaled consumption CAPM.

Only the scaled consumption CAPM specification is able to eliminate residual book-to-market effects. For this model, the coefficient on the included variable book-market ratio is not statistically different from zero at conventional levels when either the uncorrected or Shanken-corrected standard errors are used. The conditional CCAPM is the only CAPM model that passes this test. In addition, a comparison with the results in table 3 shows that there is no substantial increase in adjusted R^2 from including the portfolio book-market ratio in the scaled consumption CAPM specifications. Thus the results presented in this subsection provide no evidence that misspecification bias due to the inclusion of useless factors is driving the earlier findings for the scaled consumption CAPM.

In summary, the scaled consumption CAPM performs better in explaining the cross section of returns than the other scaled models along at least two dimensions. First, in contrast to the scaled human capital CAPM, the average risk price for the consumption beta has the correct sign. Second, in contrast to all the other scaled models, portfolio characteristics do not show up as significant explanatory variables once the scaled factors are included. For these reasons, the scaled CCAPM is our preferred specification.

F. *Alternative Scaling Variables*

So far we have considered scaled multifactor models using \widehat{cay} as a conditioning variable. We argued above that \widehat{cay} is an excellent candidate for a scaling variable and may have important advantages over other popular conditioning variables because both theory and empirical evidence suggest that it summarizes investor expectations about future returns on the entire market portfolio rather than merely some component of the market portfolio return. Still, we also compared the results presented in this paper with those using the same empirical technique and alternative scaling variables, such as the log dividend-price ratio, the spread between yields on a BAA-rated bond and a AAA-rated bond (Jagannathan and Wang 1996), and a term spread. These results are available on request. With few exceptions, these results show that a conditional CCAPM, using \widehat{cay} as conditioning information, does a significantly better job of explaining the cross section of average returns than conditional scaled models in which these alternative indicators are used as instruments.

G. *Scaled Returns*

Our chief concern in this paper is not to test the scaled versions of the (C)CAPM on every set of portfolios available, but rather to explain the cross section of returns on those 25 portfolios that have so bedeviled the CAPM and so lauded the Fama-French three-factor model. Nevertheless, in this subsection, we briefly discuss estimation results using *scaled returns*, which is analogous to Hansen and Singleton's use of conditioning variables in generalized method of moments estimation. As Cochrane (1996) emphasizes, scaled returns are intuitively appealing because they may be interpreted as "managed" returns, whereby a manager invests more or less in the unscaled portfolios according to the signal provided by the scaling variable, in our case the lagged value of \widehat{cay} .³⁰

We perform two additional Fama-MacBeth regressions including scaled portfolio returns. First, we add two returns, the market return, $R_{m,t}$ and the scaled market return, $R_{m,t}\widehat{cay}_{t-1}$, to the set of 25 size and book-to-market returns to obtain a new cross section of 27 portfolio returns. Second, we multiply each of the 25 returns by \widehat{cay}_{t-1} and add these to the original 25 returns, for a total of 50 returns. The results of

³⁰ Cochrane (1996) also emphasizes that it is important to scale the scaling variable so that the moments of the scaled returns are roughly comparable to those of the unscaled returns. Otherwise, scaled returns can have unrealistic units. Thus we follow Cochrane and use $1 + [\widehat{cay}_{t-1}/\sigma(\widehat{cay})]$ to scale returns, where $\sigma(\widehat{cay})$ is the standard deviation of \widehat{cay} . Recall that \widehat{cay}_{t-1} is demeaned, so this just multiplies each return by the standardized value of \widehat{cay}_{t-1} while preserving the scale of the return.

TABLE 7
SCALED RETURNS

Portfolios	Joint Significance (1)	R^2 (2)	\bar{R}^2 (3)	Pricing Error (4)
A. CCAPM				
25 Fama-French, market return, and scaled market return	.81	.06	.05	.77
25 Fama-French and scaled 25 Fama-French	.11	.14	.12	1.57
B. Scaled CCAPM				
25 Fama-French, market return, and scaled market return	.00	.71	.68	.41
25 Fama-French and scaled 25 Fama-French	.00	.71	.69	.91
C. Fama-French				
25 Fama-French, market return, and scaled market return	.00	.58	.53	.50
25 Fama-French and scaled 25 Fama-French	.01	.46	.43	1.24

NOTE.—See the notes to tables 1–3. This table presents results of Fama-MacBeth regressions using scaled portfolio returns in addition to the 25 Fama-French portfolios. Results reported in the top row of each panel are for regressions in which the market return and the scaled market return are added to the original 25 portfolio returns. Results reported in the bottom row of each panel are for regressions in which the scaled returns of all 25 Fama-French portfolios are added to the original 25 portfolio returns. The scaled portfolio returns are computed as $(1 + [\widehat{\alpha}_i] / \sigma(\widehat{\alpha}_i)) R_{i,t+1}$. Col. 1 reports the joint significance of all factors, col. 2 reports the unadjusted R^2 statistic for the cross-sectional regression, col. 3 reports the R^2 statistic adjusted for degrees of freedom, and col. 4 reports the square root of the average squared pricing error across all portfolios.

the Fama-MacBeth cross-sectional tests for these two sets of portfolios are presented in table 7. The table reports, for the unscaled CCAPM, the scaled CCAPM, and the Fama-French three-factor model, the adjusted and unadjusted R^2 statistics for the cross-sectional regression and the square root of the average squared pricing errors. To avoid clutter, this table reports only the p -values for joint significance of all the factors in each model rather than t -statistics for each individual factor.

Table 7 shows that the unscaled CCAPM does a poor job of explaining both sets of portfolios, which include scaled returns. For the set of 27 portfolios, the adjusted R^2 statistic is only 5 percent; for the set of 50 portfolios, it is only 12 percent. By contrast, the scaled CCAPM generates an adjusted R^2 statistic for these portfolios of 68 and 69 percent, respectively, and average pricing errors that are much smaller than those for the unscaled CAPM. Indeed, they are about half as large when the set of 27 returns is used and about 60 percent as large when the set of 50 returns is used. Interestingly, the scaled CCAPM also does better than the Fama-French model in capturing the cross section of returns on these two sets of portfolios, which include scaled returns. The Fama-French model explains about 53 percent of the cross-sectional variation in the 27 returns, compared to 68 percent for the scaled CCAPM, and

the square root of the average squared pricing errors for the Fama-French model is about 20 percent larger than that of the scaled CCAPM. The Fama-French model explains about 43 percent of the cross-sectional variation on the 50 portfolio returns, compared to 69 percent for the scaled CCAPM, and the square root of the average squared pricing errors for the Fama-French model is almost 40 percent larger than that of the scaled CCAPM.

These results imply that the Fama-French model has greater difficulty pricing the scaled portfolio returns than the scaled CCAPM model has. Note that the cross-sectional results using scaled returns can differ from those using the unscaled results only if the scaling variable has some time-series forecasting power for the returns being scaled (i.e., if the covariance between $z_{i,t-1}$ and $R_{i,t}$ is nonzero). We find that $\widehat{cay}_{i,t-1}$ (unlike other popular forecasting variables, e.g., the dividend yield) does in fact have forecasting power for these returns, with R^2 statistics from quarterly forecasting regressions ranging from 9 to 16 percent for the 25 portfolios.

V. Alternative Estimation Methodologies

In this paper, we use the Fama-MacBeth procedure for testing the cross-sectional explanatory power of each model. As Cochrane (2001) notes, the Fama-MacBeth methodology is practically the same as first-stage generalized method of moments (GMM), where the identity weighting matrix is used. Other methodologies exist, including some form of second-stage GMM (where an estimated weighting matrix is used).³¹ This type of second-stage GMM, using the second-moment matrix for returns as a weighting matrix, has been advocated by Roll and Ross (1994) and Kandel and Stambaugh (1995) as a way of checking whether cross-sectional asset pricing tests are sensitive to the particular set of portfolios on which the tests were carried out. For the application pursued in this paper, however, we argue that the Fama-MacBeth methodology—or first-stage GMM—is more appropriate than second-stage GMM, where an estimated weighting matrix is employed. We make this argument for two reasons.

First, second-stage GMM estimation is unsuitable in studies that have a small time-series sample relative to the cross-sectional sample size (Ferson and Foerster 1994; Altonji and Segal 1996; Christiano and Den Haan 1996; Hansen et al. 1996; Ahn and Gadarowski 1999). Our quarterly data yield far fewer time-series observations than in many previous

³¹ We use the term “second-stage” loosely here to refer to any GMM estimation in which the matrix used to weight the criterion function is estimated, even if the GMM procedure can be done in one “stage.”

empirical applications for which monthly data were available. Both the optimal GMM weighting matrix of Hansen (1982) and the second-moment matrix of returns (Hansen and Jagannathan 1997) are likely to be poorly estimated in samples of the size encountered here. Indeed, for this reason, it is often argued that first-stage GMM should serve as a robustness check on the results of any second-stage GMM estimation.³² If the two sets of results differ greatly, the studies cited above suggest that the source of discrepancy may lie with the poor finite sample estimate of the weighting matrix used in second-stage GMM.

To illustrate the potential for this problem in our application, we performed GMM estimation of our empirical specification for M_{t+1} on the full set of 25 portfolios using the second-moment matrix of returns to weight the criterion function and computed the Hansen and Jagannathan (1997) distance measure along with its associated p -value as in Jagannathan and Wang (1996).³³ Given the small-sample problems with GMM, it is not surprising that this estimation produces a large Hansen-Jagannathan distance that is statistically different from zero for *all* the models we considered (including the Fama-French model). These statistical criteria apparently have little power to discriminate among the asset pricing models we explore, suggesting that the Fama-French model does about the same as the unconditional CCAPM, which does about the same as the scaled (C)CAPM models in explaining the cross section of average returns on these portfolios. Hodrick and Zhang (2000) report similar findings using quarterly data of roughly the same sample used here. To check whether these findings are likely to be a mere artifact of our small time-series sample (which contains 146 quarterly observations) given the relatively large cross section of 25 portfolios, we perform three applications of GMM estimation that, according to the research cited above, should be more robust to small-sample biases.

First, as recommended by Altonji and Segal (1996), we undertake GMM estimation on the full set of 25 Fama-French portfolios, but instead of using the second-moment matrix of returns to weight the criterion function, we use the identity matrix. It should not be surprising that these results lead to the same conclusions as the Fama-MacBeth pro-

³² Altonji and Segal (1996) show that first-stage GMM estimates using the identity matrix are far more robust to small-sample problems than GMM estimates in which the criterion function has been weighted with an estimated matrix. Cochrane (2001) recommends using the identity matrix as a robustness check in any estimation in which the cross-sectional dimension of the sample is less than one-tenth of the time-series dimension.

³³ See Hansen and Jagannathan (1997) and Jagannathan and Wang (1996) for a detailed explanation of the Hansen-Jagannathan distance. It gives the squared distance from the candidate stochastic discount factor of a given asset pricing model to the set of all the discount factors that price the N assets correctly. Alternatively, the Hansen-Jagannathan distance is the pricing error for the portfolio that is most mispriced by the model. In principle, if the model is correct, the Hansen-Jagannathan distance should be statistically indistinguishable from zero.

cedure, since, as mentioned, the two approaches are essentially the same. By explicitly running first-stage GMM, however, we may readily compute an alternative statistical measure of model performance, namely the square root of the minimized GMM objective function, a statistic we denote *Dist*. Following Jagannathan and Wang (1996)—who show how one can test whether *Dist* is statistically different from zero for any GMM estimation using an arbitrary weighting matrix—we test whether this distance measure is statistically different from zero. This provides a formal test of each asset pricing model.

Second, an alternative way to reduce small-sample bias is to cut back on the number of cross-sectional observations by performing second-stage GMM estimation on a smaller set of portfolios. We do so for a subset of 10 of the original 25 Fama-French portfolios and for a set of six size and book-market sorted portfolios provided by Fama and French. For both of these second-stage estimations, we use the second-moment matrix of returns as a weighting matrix. These last two GMM estimations minimize an objective function that is insensitive to the initial choice of portfolios (so it addresses the criticisms raised in Roll and Ross [1994] and Kandel and Stambaugh [1995]) but at the same time reduce the number of cross-sectional observations (thereby mitigating the small-sample biases).

The results are presented in table 8, and the methodology is discussed in more detail in Appendix B. For each of these three alternative GMM applications (one using the full set of 25 portfolios and the identity weighting matrix and two using a reduced number of portfolios and the second-moment weighting matrix), the scaled consumption CAPM, with *cay* used as a scaling variable, posts a significant improvement over its unscaled CCAPM counterpart. For all three of these estimations, the statistic *Dist* is between 47 and 63 percent smaller for the scaled CCAPM than it is for the unscaled CCAPM and close in magnitude to the *Dist* measure for the Fama-French three-factor model. Moreover, we cannot reject the hypothesis that $\text{Dist} = 0$ for the scaled CCAPM, whereas the opposite is true for the unscaled CCAPM. These results contrast sharply with those when the full set of 25 portfolios are used, and they reject the conclusion drawn in Hodrick and Zhang (2000) that scaling does little to enhance the performance of the CCAPM when these types of tests are performed. When we take steps to reduce the potential impact of small-sample biases, scaling with \widehat{cay} improves model performance by both an economically and statistically significant amount. These findings underscore the danger of running empirical horse races among asset pricing models by performing GMM estimation in small samples in which the weighting matrix is estimated.

Table 8 also shows that we can reject the Fama-French three-factor model (*Dist* is statistically different from zero), even though its pricing

TABLE 8
GMM ESTIMATES

PORTFOLIOS AND WEIGHTING MATRICES	CCAPM		SCALED CCAPM		FAMA-FRENCH	
	Joint Significance (1)	Dist (2)	Joint Significance (3)	Dist (4)	Joint Significance (5)	Dist (6)
25:						
Hansen-Jagannathan matrix	.47 (.00)	65.44 (.00)	.01 (.00)	62.75 (.00)	.00 (.00)	59.33 (.00)
Identity matrix	.27 (.01)	3.14 (.01)	.00 (.01)	1.86 (.14)	.00 (.01)	1.54 (.00)
10:						
Hansen-Jagannathan matrix	.54 (.00)	47.92 (.00)	.00 (.00)	38.46 (.24)	.00 (.00)	40.60 (.00)
Identity matrix	.36 (.01)	2.00 (.01)	.00 (.01)	1.26 (.20)	.00 (.01)	1.12 (.00)
6:						
Hansen-Jagannathan matrix	.23 (.01)	38.07 (.01)	.00 (.01)	26.03 (.14)	.00 (.01)	28.99 (.01)
Identity matrix	.25 (.09)	1.21 (.09)	.00 (.09)	.57 (.10)	.00 (.09)	.57 (.01)

NOTE.—See the notes for tables 1–3. This table presents GMM estimates of models of the form

$$1 = E[M_{t+1}(1 + R_{t+1})], \quad M_{t+1} = \mathbf{b}'\mathbf{F}_{t+1}.$$

The consumption CAPM uses only consumption growth Δc as factor, and the scaled C-CAPM uses \widehat{cay} as a scaling variable for Δc . The three-factor Fama-French model includes the value-weighted CRSP index (R_m) and the Fama-French factors SMB and HML. The model is estimated using three sets of portfolios: the original 25 Fama-French portfolios, a subset of 10 portfolios (S1B1, S1B5, S2B1, S1B5, S3B1, S3B5, S4B1, S4B5, S5B1, and S5B5), and the six portfolios underlying SMB and HML. Two different weighting matrices are used: the second-moment matrix of returns (Hansen and Jagannathan 1991) and the identity matrix. Cols. 1, 3, and 5 report p -values of a test of joint significance of the estimated parameters in \mathbf{b} . Dist denotes the square root of the minimized GMM objective function (multiplied by 100). p -values of a test that Dist is equal to zero are reported in parentheses (computed according to Jagannathan and Wang [1996]). The model is estimated using data from 1963:Q3 to 1998:Q3.

errors are roughly the same as those of scaled CCAPM. Like the Wald test results discussed above, these findings point to a potential problem with using purely statistical criteria to judge model performance, particularly in small samples in which second moments are likely to be poorly estimated.³⁴

Finally, there is a more fundamental reason we favor the Fama-

³⁴ Other researchers have recommended the use of generalized least squares (GLS) on the grounds that asset returns may display conditional heteroskedasticity. When conditional heteroskedasticity is present, the GLS approach should improve efficiency. However, the problems with GLS estimation are similar to those with second-stage GMM estimation. First, in small samples, the GLS transformation can place too much weight on what appear to be nearly riskless portfolios so measured as a result of luck in a short sample. Second, as in any application of GLS, the improvement in efficiency depends on knowing the true covariance matrix of returns. Since this knowledge is rare, GLS is often less robust than the Fama-MacBeth procedure based on OLS. Third, conditional heteroskedasticity in quarterly data is less evident than in the monthly data commonly used, so the improvement in efficiency from GLS is likely to be marginal. Even so, we may obtain parameter estimates without making the assumption that the errors are conditional homoskedastic by using first-stage GMM to estimate \mathbf{b} and then using (5) to back out the λ . Appendix B demonstrates that doing so gives values for λ that are almost identical to those using the Fama-MacBeth methodology.

MacBeth approach over second-stage GMM of any kind. The original Fama-French portfolios, sorted according to size and book-to-market equity ratios, were chosen carefully to represent economically interesting characteristics. When an estimated weighting matrix is used, test portfolios become linear combinations of the original portfolios, which can be difficult to interpret economically and can even imply implausible long and short positions in the original assets. The size and value puzzles documented by Fama and French (1992, 1993) originate in reference to the original 25 portfolios, not to some reweighted portfolio of these portfolios.³⁵ Thus, if one wants to address the question of whether a set of macroeconomic factors can explain the value premium documented by Fama and French (1992), it is necessary to focus asset pricing tests on the original size and book-market sorted portfolios. Use of the second-moment weighting matrix undoes this specification of carefully constructed portfolios based on economically interesting characteristics.

VI. Conclusion

Empirical asset pricing has presented an abundance of formidable challenges for both the CAPM and the consumption CAPM in recent years. One of the most compelling of these challenges is presented in two papers by Fama and French (1992, 1993), who show that a broad stock market beta cannot explain the difference in return between portfolios with high and low book-to-market equity ratios. Consumption-based asset pricing models fare little better in this regard.

The failures of the CAPM and the consumption CAPM documented over the last 15 years have prompted researchers to seek alternative empirical models for explaining the pattern of returns on portfolios formed according to size and book-to-market equity ratios. Fama and French (1993) demonstrate that a three-factor model consisting of a broad stock market beta and betas on two mimicking portfolios related to size and book-to-market equity ratios can capture strong common variation in returns. Yet these results have been a source of controversy as some researchers question whether the mimicking portfolios truly capture nondiversifiable, and therefore macroeconomic, risk. Since models that specify actual macroeconomic variables as risk factors have, as yet, failed to explain a significant fraction of the variation in these returns, this contention persists.

We argue that the results presented in this paper go a long way toward

³⁵ This is evidenced by the finding that the Fama-French three-factor model—designed to fit the empirical evidence on the original portfolios—does poorly in GMM tests in which the second-moment matrix of returns is used to weight the criterion function.

resolving this controversy. We employ an empirical test of the (C)CAPM in which the discount factor is approximated as a linear function of the model's fundamental factors. Instead of assuming that the parameters of this function are fixed over time, as in many previous studies, we model the parameters as time-varying by scaling them with a proxy for the log consumption–wealth ratio. In contrast to the simple static CAPM or unconditional consumption CAPM, we find that these scaled multifactor versions of the CCAPM can explain a substantial fraction of the cross-sectional variation in average returns on stock portfolios sorted according to size and book-to-market equity ratios. These results seem to be especially supportive of a habit formation version of the consumption CAPM, where the multiplicative, or scaled, consumption factor is important. This scaled consumption CAPM does a good job of explaining the celebrated value premium: portfolios with high book-to-market equity ratios also have returns that are more highly correlated with the scaled consumption factors we consider, and vice versa. Furthermore, the scaled consumption model eliminates residual size and book-to-market effects that remain in the CAPM. Taken together, these findings lend support to the view that the value premium is at least partially attributable to the greater nondiversifiable risk of high-book-to-market portfolios, and not simply to elements bearing no relation to risk, such as firm characteristics or sample selection biases.

Our results also help shed light on *why* the Fama-French three-factor model performs so well relative to the unscaled size: the data suggest that the Fama-French factors are mimicking portfolios for risk factors associated with time variation in risk premia. Once the (C)CAPM is modified to account for such time variation, it performs about as well as the Fama-French model in explaining the cross-sectional variation in average returns. Of course, as with any model, the one investigated here is only an approximation of reality, and it is clear that some features of these data remain unexplained even after one accounts for these consumption covariances. The success of the (C)CAPM model tested here rests with its relative accuracy rather than with its ability to furnish a flawless description of reality.

A key component of this success is our choice of conditioning information. We argue here that the difference between log consumption and a weighted average of log asset wealth and log labor income is likely to provide a superior summary measure of conditional expectations. We find that, consistent with this proposition, the scaled consumption CAPM, using \widehat{cay} as an instrument, typically performs far better than it does using other possible instruments, such as the dividend-price ratio, the default spread, or the term spread.

The conditional linear factor models we explore here are quite different from unconditional models. If conditional expected returns to

the market portfolio are time-varying, the investor's discount factor will not merely depend unconditionally on consumption growth or the market return, but instead will be a function of these factors conditional on information about future returns. Assets are riskier if their returns are more highly *conditionally* correlated with factors, rather than unconditionally correlated as in classic versions of the (C)CAPM. The approach taken here, of scaling factors with information available in the current period, leads to a multifactor, unconditional model in place of a single-factor, conditional model. This approach therefore provides a justification for requiring more than one factor to explain the behavior of expected returns, even if one believes that the true model is, for example, a consumption-based intertemporal asset pricing model with a single fundamental factor. By deriving this multifactor structure from an equilibrium framework, we can mitigate a common criticism of multifactor models, namely, that multiple factors are chosen without regard to economic theory.³⁶ The empirical results we obtain from doing so suggest that a multifactor version of the consumption CAPM can explain a large portion of the cross section of expected stock returns.

Appendix A

Derivation of the Approximate Log Consumption–Wealth Ratio

The approximation of the log consumption–aggregate wealth ratio presented here was first derived in Campbell and Mankiw (1989). They show that the investor's intertemporal budget constraint, $W_{t+1} = (1 + R_{m,t+1})(W_t - C_t)$, may be expressed as

$$\Delta w_{t+1} \approx k + r_{m,t+1} + \left(1 - \frac{1}{\rho_w}\right)(c_t - w_t), \quad (\text{A1})$$

where W_{t+1} is aggregate (human plus nonhuman) wealth in period $t + 1$; ρ_w is the steady-state ratio of invested to total wealth, $(W - C)/W$; and k is a linearization constant that plays no role in our analysis. Solving this difference equation forward and imposing that $\lim_{i \rightarrow \infty} \rho_w^i (c_{t+i} - w_{t+i}) = 0$, we may write the log consumption–wealth ratio as

$$c_t - w_t = \sum_{i=1}^{\infty} \rho_w^i (r_{m,t+i} - \Delta c_{t+i}). \quad (\text{A2})$$

Taking expectations of (A2) produces (9).

³⁶ For example, these criticisms can be found in Lo and MacKinlay (1990), Breen and Korajczyk (1993), and Kothari, Shanken, and Sloan (1995).

Appendix B

GMM Estimation

This Appendix presents results from generalized method of moments (GMM) estimates of models taking the form

$$1 = E_t[M_{t+1}(1 + R_{i,t+1})], \quad M_{t+1} = \mathbf{b}'\mathbf{F}_{t+1}.$$

We use two prespecified weighting matrices, the Hansen-Jagannathan, second-moment weighting matrix and the identity matrix. To assess whether the GMM results may be influenced by small-sample problems, three alternative GMM estimations are performed. Following the suggestion of Altonji and Segal (1996), we present first-stage GMM estimates using the identity weighting matrix. We also performed GMM estimation on a subset of 10 of the original 25 Fama-French portfolios, using a set of six size and book-market sorted portfolios provided by Fama and French. These six portfolios consist of three size and two book-market categories of the same data used to construct the five size and five book-market portfolios. Table 8 in the text presents a distance measure for each GMM estimation, Dist (the square root of the minimized GMM objective function), as well as the joint significance level of all elements in the parameter vector \mathbf{b} , for each model. The variable Dist is a summary statistic that can be used to compare the size of the pricing errors across models; it gives the squared distance from the candidate stochastic discount factor and the set of discount factors that price the 25 portfolios correctly. If the model is correct, Dist should equal zero.

The table shows that scaling the standard CCAPM does not decrease Dist significantly when the weighting matrix is the Hansen-Jagannathan matrix and the estimation is run on the full set of 25 portfolios. Furthermore, results (not reported) showed that this estimation rejects the hypothesis that Dist = 0 for all models, including the Fama-French model. By contrast, when the weighting matrix is the identity matrix, the effect of scaling the consumption CAPM is to reduce Dist by almost half, and we cannot reject the null hypothesis that the distance is zero for the scaled consumption CAPM. Qualitatively identical results are produced when we use the Hansen-Jagannathan matrix, but the number of portfolios is reduced, either by using a subset of the original 25 portfolios or by using the six portfolios. In each case, scaling the CCAPM significantly improves the model's performance, and we do not reject the hypothesis that the pricing errors are zero. Also, the parameters in \mathbf{b} are strongly statistically significant for the scaled CCAPM.

The estimated coefficients in \mathbf{b} for the scaled consumption CAPM in the case in which estimation was carried out on the full 25 portfolios using the identity weighting matrix are $\hat{b}_0 = 0.99$, $\hat{b}_{\Delta c} = -26.47$, $\hat{b}_{\Delta s} = -16,076$, and $\hat{b}_z = 76.71$. To check whether the GMM estimation produces results that are similar to those produced by the Fama-MacBeth methodology, we can convert the b 's into λ 's using (5). We obtain for this case (using the same scale as in table 4) $E[R_{0,t}] = 4.31$, $\lambda_{\Delta c} = 0.02$, $\lambda_{\Delta s} = 0.06$, and $\lambda_z = -0.14$. Note that these numbers are extremely close to those in row 2 of table 3. Note that estimates of \mathbf{b} are obtained without making the assumption that the errors are conditionally homoskedastic, as is the case for the Shanken-corrected Fama-MacBeth standard errors. See Jagannathan and Wang (1998) for a derivation of the Fama-MacBeth standard errors when conditional heteroskedasticity is present.

References

- Ahn, Seung C., and Gadarowski, Christopher. "Small Sample Properties of the Model Specification Test Based on the Hansen-Jagannathan Distance." Manuscript. Tempe: Arizona State Univ., 1999.
- Altonji, Joseph G., and Segal, Lewis M. "Small-Sample Bias in GMM Estimation of Covariance Structures." *J. Bus. and Econ. Statist.* 14 (July 1996): 353–66.
- Avramov, Doron. "Stock-Return Predictability and Model Uncertainty." Working paper. Philadelphia: Univ. Pennsylvania, Wharton School, 1999.
- Banz, Rolf W. "The Relationship between Return and Market Value of Common Stocks." *J. Financial Econ.* 9 (March 1981): 3–18.
- Basu, S. "The Investment Performance of Common Stocks in Relation to Their Price-Earnings Ratios: A Test of the Efficient Market Hypothesis." *J. Finance* 32 (June 1977): 663–82.
- Berk, Jonathan B. "A Critique of Size-Related Anomalies." *Rev. Financial Studies* 8 (Summer 1995): 275–86.
- Black, Fischer. "Capital Market Equilibrium with Restricted Borrowing." *J. Bus.* 45 (July 1972): 444–55.
- Black, Fischer; Jensen, Michael C.; and Scholes, Myron S. "The Capital Asset Pricing Model: Some Empirical Tests." In *Studies in the Theory of Capital Markets*, edited by Michael C. Jensen. New York: Praeger, 1972.
- Blume, Marshall E., and Friend, Irwin. "A New Look at the Capital Asset Pricing Model." *J. Finance* 28 (March 1973): 19–33.
- Breeden, Douglas T. "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities." *J. Financial Econ.* 7 (September 1979): 265–96.
- Breeden, Douglas T.; Gibbons, Michael R.; and Litzenberger, Robert H. "Empirical Tests of the Consumption-Oriented CAPM." *J. Finance* 44 (June 1989): 231–62.
- Breeden, Douglas T., and Litzenberger, Robert H. "Prices of State-Contingent Claims Implicit in Option Prices." *J. Bus.* 51 (October 1978): 621–51.
- Breen, William, and Korajczyk, Robert. "On Selection Biases in Book-to-Market Based Tests of Asset Pricing Models." Working Paper no. 167. Evanston, Ill.: Northwestern Univ., 1993.
- Burnside, Craig, and Eichenbaum, Martin S. "Small-Sample Properties of GMM-Based Wald Tests." *J. Bus. and Econ. Statist.* 14 (July 1996): 294–308.
- Campbell, John Y. "A Variance Decomposition for Stock Returns." *Econ. J.* 101 (March 1991): 157–79.
- . "Understanding Risk and Return." *J.P.E.* 104 (April 1996): 298–345.
- Campbell, John Y., and Cochrane, John H. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *J.P.E.* 107 (April 1999): 205–51.
- . "Explaining the Poor Performance of Consumption-Based Asset Pricing Models." *J. Finance* 55 (December 2000): 2863–78.
- Campbell, John Y., and Mankiw, N. Gregory. "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence." In *NBER Macroeconomics Annual*, edited by Olivier J. Blanchard and Stanley Fischer. Cambridge, Mass.: MIT Press, 1989.
- Campbell, John Y., and Shiller, Robert J. "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors." *Rev. Financial Studies* 1 (Fall 1988): 195–228.
- Chang, Ganlin, and Sundaresan, Suresh M. "Asset Prices and Default-Free Term

- Structure in an Equilibrium Model of Default." Working paper. New York: Columbia Univ., 1999.
- Christiano, Lawrence J., and Den Haan, Wouter J. "Small-Sample Properties of GMM for Business-Cycle Analysis." *J. Bus. and Econ. Statist.* 14 (July 1996): 309–27.
- Christoffersen, Peter F., and Diebold, Francis X. "How Relevant Is Volatility Forecasting for Financial Risk Management?" *Rev. Econ. and Statist.* 82 (February 2000): 12–22.
- Cochrane, John H. "A Cross-Sectional Test of an Investment-Based Asset Pricing Model." *J.P.E.* 104 (June 1996): 572–621.
- . *Asset Pricing*. Princeton, N.J.: Princeton Univ. Press, 2001.
- Constantinides, George M. "Habit Formation: A Resolution of the Equity Premium Puzzle." *J.P.E.* 98 (June 1990): 519–43.
- Constantinides, George M., and Duffie, Darrell. "Asset Pricing with Heterogeneous Consumers." *J.P.E.* 104 (April 1996): 219–40.
- Daniel, Kent, and Titman, Sheridan. "Evidence on the Characteristics of Cross Sectional Variation in Stock Returns." *J. Finance* 52 (March 1997): 1–33.
- Davis, James L.; Fama, Eugene F.; and French, Kenneth R. "Characteristics, Covariances, and Average Returns: 1929 to 1997." *J. Finance* 55 (February 2000): 389–406.
- Elton, Edwin J.; Gruber, Martin J.; and Blake, Christopher R. "Fundamental Economic Variables, Expected Returns, and Bond Fund Performance." *J. Finance* 50 (September 1995): 1229–56.
- Fama, Eugene F., and French, Kenneth R. "Dividend Yields and Expected Stock Returns." *J. Financial Econ.* 22 (October 1988): 3–25.
- . "Business Conditions and Expected Returns on Stocks and Bonds." *J. Financial Econ.* 25 (November 1989): 23–49.
- . "The Cross-Section of Expected Returns." *J. Finance* 47 (June 1992): 427–65.
- . "Common Risk Factors in the Returns on Stocks and Bonds." *J. Financial Econ.* 33 (February 1993): 3–56.
- . "Size and Book-to-Market Factors in Earnings and Returns." *J. Finance* 50 (March 1995): 131–55.
- . "Multifactor Explanations of Asset Pricing Anomalies." *J. Finance* 51 (March 1996): 55–84.
- Fama, Eugene F., and MacBeth, James D. "Risk, Return, and Equilibrium: Empirical Tests." *J.P.E.* 81 (May/June 1973): 607–36.
- Fama, Eugene F., and Schwert, G. William. "Human Capital and Capital Market Equilibrium." *J. Financial Econ.* 4 (January 1977): 95–125.
- Ferson, Wayne E., and Foerster, Stephen. "Finite Sample Properties of the Generalized Method of Moments in Tests of Conditional Asset Pricing Models." *J. Financial Econ.* 36 (August 1994): 29–55.
- Ferson, Wayne E., and Harvey, Campbell R. "Seasonality and Heteroscedasticity in Consumption-Based Asset Pricing: An Analysis of Linear Models." In *Research in Finance*, vol. 11, edited by Andrew H. Chen. Greenwich, Conn.: JAI, 1993.
- . "Conditioning Variables and the Cross Section of Stock Returns." *J. Finance* 54 (August 1999): 1325–60.
- Ferson, Wayne E.; Kandel, Shmuel; and Stambaugh, Robert F. "Tests of Asset Pricing with Time-Varying Expected Risk Premiums and Market Betas." *J. Finance* 42 (June 1987): 201–20.
- Ferson, Wayne E.; Sarkissian, Sergei; and Simin, Timothy. "The Alpha Factor Asset Pricing Model: A Parable." *J. Financial Markets* 2 (February 1999): 49–68.

- . “Spurious Regressions in Financial Economics?” Working paper. Seattle: Univ. Washington, 2000.
- Hansen, Lars Peter. “Large Sample Properties of Generalized Method of Moments Estimators.” *Econometrica* 50 (July 1982): 1029–54.
- Hansen, Lars Peter; Heaton, John; and Yaron, Amir. “Finite-Sample Properties of Some Alternative GMM Estimators.” *J. Bus. and Econ. Statist.* 14 (July 1996): 262–80.
- Hansen, Lars Peter, and Jagannathan, Ravi. “Assessing Specification Errors in Stochastic Discount Factor Models.” *J. Finance* 52 (June 1997): 557–90.
- Hansen, Lars Peter, and Singleton, Kenneth J. “Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models.” *Econometrica* 55 (September 1982): 1269–86.
- . “Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns.” *J.P.E.* 91 (April 1983): 249–65.
- Harrison, J. Michael, and Kreps, David M. “Martingales and Arbitrage in Multiperiod Securities Markets.” *J. Econ. Theory* 20 (June 1979): 381–408.
- Harvey, Campbell R. “Time-Varying Conditional Covariances in Tests of Asset Pricing Models.” *J. Financial Econ.* 24 (October 1989): 289–317.
- Hodrick, Robert J. “Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement.” *Rev. Financial Studies* 5, no. 3 (1992): 357–86.
- Hodrick, Robert J., and Zhang, Xiaoyan. “Evaluating the Specification Errors of Asset Pricing Models.” Manuscript. New York: Columbia Univ., 2000.
- Jagannathan, Ravi, and Wang, Zhenyu. “The Conditional CAPM and the Cross-Section of Expected Returns.” *J. Finance* 51 (March 1996): 3–53.
- . “An Asymptotic Theory for Estimating Beta-Pricing Models Using Cross-Sectional Regression.” *J. Finance* 53 (August 1998): 1285–1309.
- Kan, Raymond, and Zhang, Chu. “Two-Pass Tests of Asset Pricing Models with Useless Factors.” *J. Finance* 54 (February 1999): 203–35.
- Kandel, Shmuel, and Stambaugh, Robert F. “Portfolio Inefficiency and the Cross-Section of Expected Returns.” *J. Finance* 50 (March 1995): 157–84.
- Kothari, S. P.; Shanken, Jay; and Sloan, Richard G. “Another Look at the Cross-Section of Expected Stock Returns.” *J. Finance* 50 (March 1995): 185–224.
- Lakonishok, Josef; Shleifer, Andrei; and Vishny, Robert W. “Contrarian Investment, Extrapolation, and Risk.” *J. Finance* 49 (December 1994): 1541–78.
- Lamont, Owen. “Earnings and Expected Returns.” *J. Finance* 53 (October 1998): 1563–87.
- Lettau, Martin, and Ludvigson, Sydney. “Consumption, Aggregate Wealth, and Expected Stock Returns.” *J. Finance* 56 (June 2001): 815–49.
- Liew, Jimmy, and Vassalou, Maria. “Can Book-to-Market, Size and Momentum Be Risk Factors That Predict Economic Growth?” *J. Financial Econ.* 57 (August 2000): 221–45.
- Lintner, John. “Security Prices, Risk, and Maximal Gains from Diversification.” *J. Finance* 20 (December 1965): 587–615.
- Lo, Andrew, and MacKinlay, A. Craig. “Data-Snooping Biases in Tests of Financial Asset Pricing Models.” *Rev. Financial Studies* 3, no. 3 (1990): 431–67.
- Mankiw, N. Gregory, and Shapiro, Matthew D. “Risk and Return: Consumption Beta versus Market Beta.” *Rev. Econ. and Statist.* 68 (August 1986): 452–59.
- Mayers, David. “Nonmarketable Assets and Capital Market Equilibrium under Uncertainty.” In *Studies in the Theory of Capital Markets*, edited by Michael C. Jensen. New York: Praeger, 1972.

- Merton, Robert C. "An Intertemporal Capital Asset Pricing Model." *Econometrica* 41 (September 1973): 867–86.
- Miller, Merton H., and Scholes, Myron S. "Rates of Return in Relation to Risk: A Re-examination of Some Recent Findings." In *Studies in the Theory of Capital Markets*, edited by Michael C. Jensen. New York: Praeger, 1972.
- Roll, Richard. "A Critique of the Asset Pricing Theory's Test: Part 1: On Past and Potential Testability of the Theory." *J. Financial Econ.* 4 (March 1977): 129–76.
- Roll, Richard, and Ross, Stephen A. "On the Cross-Sectional Relation between Expected Returns and Betas." *J. Finance* 49 (March 1994): 101–21.
- Rubinstein, Mark. "The Valuation of Uncertain Income Streams and the Pricing of Options." *Bell J. Econ.* 7 (Autumn 1976): 407–25.
- Shanken, Jay. "Multivariate Tests of the Zero-Beta CAPM." *J. Financial Econ.* 14 (September 1985): 327–48.
- . "Intertemporal Asset Pricing: An Empirical Investigation." *J. Econometrics* 45 (July–August 1990): 99–120.
- . "On the Estimation of Beta-Pricing Models." *Rev. Financial Studies* 5, no. 1 (1992): 1–33.
- Sharpe, William F. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *J. Finance* 19 (September 1964): 425–42.
- Shiller, Robert J. "Stock Prices and Social Dynamics." *Brookings Papers Econ. Activity*, no. 2 (1984), pp. 457–98.
- Stock, James H. "Asymptotic Properties of Least Squares Estimators of Cointegrating Vectors." *Econometrica* 55 (September 1987): 1035–56.
- Sundaresan, Suresh M. "Intertemporally Dependent Preferences and the Volatility of Consumption and Wealth." *Rev. Financial Studies* 2, no. 1 (1989): 73–89.